arXiv:1607.06513v2 [math.OC] 2 Aug 2016

Online First-Order Framework for Robust Convex Optimization

Nam Ho-Nguyen[1] and Fatma Kılınç-Karzan[1]

[1]Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA, 15213, USA.

August 01, 2016

Abstract

Robust optimization (RO) has emerged as one of the leading paradigms to efficiently model parameter uncertainty. The recent connections between RO and problems in statistics and machine learning domains demand for solving RO problems in ever more larger scale. However, the traditional approaches for solving RO formulations based on building and solving robust counterparts can be prohibitively expensive and thus significantly hinder the scalability of this approach. In this paper, we present a flexible iterative framework to approximately solve robust convex optimization problems. Our results are based on the recently introduced notions of weighted regret online convex optimization and online saddle point problems. In comparison to the existing literature, a key distinguishing feature of our approach is that it only requires access to cheap first-order oracles, while maintaining the same convergence rates. This in particular makes our approach much more scalable and hence preferable in large-scale applications, specifically those from machine learning and statistics domains. We also provide new interpretations of existing approaches in our framework and demonstrate our framework on two illustrative applications on support vector machines and robust quadratic programs.

1 Introduction

Robust optimization (RO) is one of the leading modeling paradigms for optimization problems under uncertainty. As opposed to the other approaches, RO seeks a solution that is immunized against *all* possible realizations of uncertain parameters (noises) from a given uncertainty set. It is widely adopted in practice mainly because of its ability to immunize problems against perturbations in model parameters while preserving computational tractability. We refer the reader to the paper by Ben-Tal and Nemirovski [6], the book by Ben-Tal et al. [4] and surveys [8, 9, 11, 14] for a detailed account of RO theory and numerous applications.

Recently, fascinating connections have been established between problems from the statistics and machine learning domains and robust optimization. More precisely, it is demonstrated that RO can be used to achieve desirable statistical properties such as stability, sparsity, and consistency. For example, for linear regression problems, El Ghaoui and Lebret [16] and Xu et al. [39] respectively establish the equivalence of the ridge regression and the Lasso to specific RO formulations of unregularized regression problems. Moreover, Xu et al. [38] exhibit similar results in the context of regularizing support vector machines (SVMs), and [38, 39] validate the statistical consistency of methods such as SVM and Lasso via RO methodology. In addition to these RO interpretations of regularization techniques used in statistics and machine learning, robust versions of many problems from these domains are gaining traction. For example, [36] examines robust variants of SVMs and

other classification problems, and [2] explores a robust formulation for kernel classification problems. We refer the reader to [14, 5] and references therein for further examples and details on connections between robust optimization and statistics and machine learning.

These recent connections not only highlight the importance of RO methodology but also present algorithmic challenges where the scalability of RO algorithms with problem dimension becomes crucial. The primary method for solving a robust convex optimization problem is to transform it into an equivalent deterministic problem called the *robust counterpart*. Under mild assumptions, this yields a convex and tractable robust counterpart problem (see [4, 11, 3]), which can then be solved using existing convex optimization software and tools. This approach has seen much success in decision making domain. A drawback of this traditional approach is that the reformulated robust counterpart is often not as scalable as the nominal program. In particular, the robust counterpart can easily belong to a different class of optimization problems as opposed to the underlying original deterministic problem. For example, a linear program (LP) with ellipsoidal uncertainty is equivalent to a convex quadratic program (QP), and similarly a conic-quadratic program with ellipsoidal uncertainty is equivalent to a semidefinite program (SDP) (see e.g., [4, 11]). Moreover, it is well-known that convex QPs as opposed to LPs, and SDPs as opposed to convex QPs are much less scalable in practice. This becomes a critical challenge in big data applications frequently encountered in machine learning and statistics, where even solving the original deterministic problem to high accuracy is prohibitively time-consuming.

The iterative schemes that alternate between the generation/update of candidate solutions and the realizations of noises offer a convenient remedy to the scalability issues associated with the robust counterpart approach. Thus far, such approaches [29] and [5] have relied on two oracles: (*i*) *solution oracles* to solve instances of extended (or nominal) problems with constraint structures similar to (or the same as) the deterministic problem, and (*ii*) *noise oracles* to generate/update particular realizations of the uncertain parameters. At each iteration of these schemes, both solution and noise oracles are called, and their outputs are used to update the inputs of each other oracle in the next iteration. Because solution oracles rely on a solver of the same class capable of solving the deterministic problem, these iterative approaches circumvent the issue of the robust counterpart approach relying on a different solver. Nevertheless, these iterative approaches still suffer from a serious drawback: the solution oracles in [29, 5] themselves can be expensive as they require solving extended or nominal optimization problems completely. While solving the nominal problem is not as computationally demanding as solving the robust counterpart, the overall procedure relying on repeated calls to such oracles can be prohibitive. In fact, each such call to a solution oracle may endure a significant computational cost, which is at least as much as the computational cost of solving an instance of the deterministic nominal problem. Note that, to ensure scalability, most applications in machine learning and statistics already need to rely on cheap first-order methods for solving deterministic nominal problems.

In this paper, we propose an efficient iterative framework for solving robust convex optimization problems which can rely on, in an *online* fashion, much cheaper *first-order oracles* in place of full solution and noise oracles. In particular, in each iteration, instead of solving a complete optimization problem within the solution and/or noise oracles, we show that simple simultaneous updates on the solution and noise in an online fashion using only first-order information from the deterministic constraint structure is sufficient to solve the RO problem. Moreover, we show that the number of calls to such online first-order (OFO) oracles is not only at most that of the approaches utilizing full optimization based oracles for solution and/or noise, but also almost independent of the

dimension of the problem. This therefore makes our approach especially attractive for applications in statistics and machine learning domains where it is critical to maintain that the overall approach has both gracious dependence on the dimension of the problem and cheap iterations. We outline our contribution more concretely after discussing the most relevant literature.

Related Work

Thus far, the iterative approaches, which bypass the restrictions of the robust counterparts, work with *extended* nominal problems that belong to the same class as the deterministic nominal one by carefully controlling the constraints included in the formulation corresponding to noise realizations.

For robust binary linear optimization problems with only objective function uncertainty and a polyhedral uncertainty set, Bertsimas and Sim [12] suggest an approach which relies on solving $n+1$ number of instances of the nominal problem, where n is the dimension of the problem.

For robust convex optimization problems, Calafiore and Campi [13] study a 'constraint sampling' approach based on forming a single extended nominal problem of the same class as the deterministic one via i.i.d. sampling of noise realizations. They show that the optimal solution to this extended nominal problem is robust feasible with high probability where the probability depends on the sampling procedure, the number of samples drawn, and the dimension.

Mutapcic and Boyd [29] follow a 'cutting-plane' type approach where in each iteration a solution oracle is called to solve an extended nominal problem of the same class as the deterministic problem and a noise oracle, referred to as *pessimization oracle*, is invoked to iteratively expand and refine the extended nominal problem. Given a candidate solution, a pessimization oracle either certifies its feasibility with respect to the robust constraints or returns a new noise realization from the uncertainty set for which the solution is infeasible; then the nominal constraint associated with that particular noise realization is included in the extended problem. This process is repeated until a robust feasible solution is found or the last extended problem is found to be infeasible. In the overall procedure, the number of iterations (or calls to the pessimization oracle) can be exponential in the dimension. Despite this, [29] reports impressive computational results.

Both of the approaches from [13] and [29] pose issues for high-dimensional problems. In [13], as the dimension grows, an extended problem with linearly more nominal constraints is required to ensure the high probability guarantee on finding a good quality solution. In [29] at each iteration, a nominal constraint is added to the extended nominal problem. The theoretical bound on the number of constraints that need to be added is exponential, so the extended problem in [29] can grow to be exponentially large. Moreover, in both cases the extended nominal problem may no longer have certain favorable problem structure of the deterministic nominal problem, such as a network flow structure.

To address these issues, in particular, the issue of solving extended nominal problems that are not only larger-in-size than the deterministic problem but also may lack certain favorable problem structure of the deterministic problem, Ben-Tal et al. [5] introduce a new iterative approach to approximately solve robust feasibility problems via a *nominal feasibility oracle* and running an online learning algorithm to choose noise realizations. Given a particular noise realization, the nominal feasibility oracle solves an instance of the deterministic nominal feasibility problem obtained by simply fixing the noise to the given value. Hence, the problem solved by this oracle has the *same* number of constraints and the *same structure* as the original nominal problem; in particular its size does not grow in each iteration. This is an important distinguishing feature of this approach. The other distinguishing feature is that Ben-Tal et al. [5] replace the pessimization oracle of [29] by employing an online learning algorithm, which simply requires first-order information of

the noise from the constraint functions. Moreover, [5] provide a dimension independent bound on the number of iterations (nominal feasibility oracle calls). Because this approach is closely related to our work, we give a detailed summary of it in Section 4.3, highlight its connections to our work; in fact we show that it can be seen as a special case of our framework.

We close with a brief summary of the assumptions on the computational requirements of these methods. The constraint sampling approach of [13] requires access to a sampling procedure on the uncertainty sets as well as an oracle capable of solving the extended nominal problem. The cutting plane approach of [29] replaces the sampling procedure of [13] with a noise oracle, namely the pessimization oracle that works with the uncertainty sets but still requires the same type of optimization oracle as a solution oracle to solve the extended problems. Ben-Tal et al. [5] substitute the pessimization oracle with an online learning-based procedure, which requires merely first-order information from the constraint functions and simple projection type operations on the associated uncertainty sets, but it still relies on a solution oracle capable of solving the original nominal problem, which is essentially the same (up to log factors) as the optimization oracles in [13] and [29]. If the deterministic problem is network flow etc, a specific solver can be used in the framework of [5], but this is not possible for [13] and [29].

Summary of Our Contributions

In fact, one can view all of these iterative approaches as two iterative processes that run simultaneously and in conjunction with each other to generate/update solutions and noise realizations. This naturally leads to a dynamic game environment where in each round Player 1 chooses a solution and Player 2 chooses a realization of uncertain parameters. In this framework, the policies employed by these players in their decision making determine the nature of the final approach. In the case of [29], Player 1 considers all of the previous noise realizations when making his decision, whereas Player 2 simply reacts to the current solution when choosing the noise. In [5], Player 1 reacts to only the current noise in generating/updating the solution while Player 2 minimizes the regret associated with past solutions in choosing noise.

In this paper, we further analyze this interaction between Player 1 and Player 2, with the aim of deriving a simpler and computationally much less demanding iterative approach to solving RO problems. Our contributions can be summarized as follows.

1. We build a *general and flexible framework* for iteratively solving robust feasibility problems, and demonstrate its flexibility by describing it as a meta-template. By customizing our framework appropriately, we modify the pessimization oracle-based approach of [29] and obtain a much better bound on the number of oracle calls in [29]. We also provide a new interpretation of the nominal feasibility oracle-based approach of [5] as a special case within our framework.

2. When the original deterministic problem admits first-order oracles capable of providing gradient/subgradient information on each constraint function, we demonstrate that *online first-order* (OFO) algorithms can be used to iteratively generate/update solutions and noise realizations simultaneously in an online manner, which leads to obtaining robust feasibility/infeasibility certificates within our framework. In contrast to the approaches of [29] and [5], which rely on full nominal feasibility oracles to generate points, our OFO-based approach only requires simple update rules in each iteration and thus has much lower per-iteration cost. Besides, our noise oracle generates a realization of the noise in an online learning fashion as was done in [5], and hence it is less expensive than the pessimization oracle of [29].

3. In our framework, the number of iterations (or oracle calls) needed to obtain approximate robust solution or a robust infeasibility certificate is a function of the approximation guarantee ϵ and the complexities of the domains for the solution and the uncertainty set; in particular, our convergence rate is (almost) dimension independent. We also demonstrate that the iteration complexity of our OFO-based approach is at least as good as that of the efficient approach of [5], and better than the exponential complexity of [29]. Overall, our OFO-based approach leads to computational savings over the approach of [5] by a factor as large as $O(1/(\epsilon^2 \log(1/\epsilon)))$ arithmetic operations when the number updates of the solution is smaller than or equal to the number of updates of the noise realization, which is the case in many applications. For further comparisons and discussion, see Section 4.4.

4. Our framework is based on formulating the robust feasibility problem as a convex-nonconcave saddle point (SP) problem, and explicitly analyzing its structure. While convex-concave SP problems are well-studied in the literature, and many efficient first-order algorithms exist for these (see for example [32, 25, 26]), the convex-*nonconcave* SP problem is not as well-studied. To our knowledge, an explicit study of convex-nonconcave SP problems and their relation to RO has not been conducted previously; in this respect, the most closely related work [5] neither provides an explicit connection between robust feasibility and SP problems, nor analyzes their structure explicitly.

5. Our framework is more amenable to exploiting favorable structural properties of the constraint functions such as strong concavity, smoothness, etc., and achieving better convergence rates. For example, we are able to employ a new weighted regret online convex optimization algorithm from [20, Section 3.2.2] for strongly convex functions with faster convergence guarantees. As a result, we show that the required number of iterations/oracle calls to obtain a robust feasible solution within the nominal feasibility oracle-based approach of [5] can be reduced by an order of magnitude; this provides a partial answer to the open question from [5] on lower bounds on the oracle calls and helps us further refine it.

In the appendix, to demonstrate the application and effectiveness of our proposed framework, we walk through two detailed examples on robust SVMs and robust QPs. In particular, for robust QPs, we are able to leverage a recent convex QP based reformulation of the classical trust region subproblem [21] in order to avoid working with a nonconvex reformulation in a lifted space as in [5, Section 4.2] and relying on a probabilistic follow-the-perturbed-leader type algorithm [5, Section 3.2]. While using such nonconvex techniques will work within our framework, our convex reformulation allows us to work directly in the original space of the variables with a deterministic subgradient-based algorithm while still achieving asymptotically similar iteration complexity guarantees as [5]. Moreover, each iteration of our approach requires only first-order updates where the most expensive operation is the computation of a maximum eigenvector; thus our per-iteration cost is significantly less.

Outline

The rest of the paper is organized as follows. We begin with some notation and preliminaries in Section 2. We introduce our robust feasibility problem and robust feasibility/infeasibility certificates in Section 2.1, convex-concave SP problems in Section 2.2, and briefly summarize important online convex optimization (OCO) tools as well as some useful OFO algorithms in Section 2.3. We formulate the robust feasibility problem as a convex-nonconcave SP problem in Section 3; this

formulation and certain bounds associated with its SP gap function form the basis of our general framework for solving robust feasibility problems. In Section 4 we specify an assortment of approaches obtained in our general framework by using different oracles. We examine our OFO-based approach in Section 4.1 by interpreting various terms in our framework in the context of OCO. In Section 4.2, we modify the pessimization oracle-based approach of [29] to obtain an efficient bound on the number of iterations required. In Section 4.3 we show how the nominal feasibility oracle-based approach of [5] fits within our framework. Finally, we discuss the convergence rates and accelerations attainable in our framework under various assumptions and compare our work with the existing approaches in Section 4.4. We close with a summary of our results and a few compelling further research directions in Section 5.

In Appendix A we give an alternative formulation of the robust feasibility problem as a convex-concave SP problem in an extended space, and discuss its andvantages and disadvantages over the convex-nonconcave SP formulation. In Appendix B we illustrate our OFO-based approach through two examples from robust SVMs and robust QPs.

2 Notation and Preliminaries

Given $a \in \mathbb{R}$, $\text{sign}(a)$ denotes the sign of the number a. For a positive integer $n \in \mathbb{N}$, we let $[n] = \{1, \ldots, n\}$ and define $\Delta_n := \{x \in \mathbb{R}^n_+ : \ \sum_{i\in[n]} x_i = 1\}$ to be the standard simplex. Throughout the paper, the superscript, e.g., f^i, u^i, U^i, is used to attribute items to the i-th constraint, whereas the subscript, e.g., x_t, f_t, ϕ_t, is used to attribute items to the t-th iteration. Therefore, we sometimes use u^i, x_t, as well as u^i_t to denote vectors in $\mathbb{R}^n$. We use the notation $\{x_t\}_{t=1}^T$ to denote the collection of items $\{x_1, \ldots, x_T\}$. Given a vector $x \in \mathbb{R}^n$, we let $x^{(k)}$ denote its k-th coordinate for $k \in [n]$. One exception we make to this notation is that we always denote the convex combination weights $\theta \in \Delta_T$ with θ_t. For $x \in \mathbb{R}^n$ and $p \in [1, \infty]$, we use $\|x\|_p$ to denote the ℓ_p-norm of x defined as

$$\|x\|_p = \begin{cases} \left(\sum_{i\in[n]} |x^{(i)}|^p\right)^{1/p} & \text{if } p \in [1, \infty) \\ \max_{i\in[n]} |x^{(i)}| & \text{if } p = \infty \end{cases}.$$

Throughout this paper, we use Matlab notation to denote vectors and matrices, i.e., $[x; y]$ denotes the concatenation of two column vectors x, y. $\mathbb{S}^n$ denotes the space of $n \times n$ symmetric matrices; and we let $\mathbb{S}^n_+$ be the positive semidefinite cone in $\mathbb{S}^n$. We let I_n denote the identity matrix in $\mathbb{S}^n$. For a matrix $A \in \mathbb{S}^n$, $\lambda_{\max}(A)$, $\|A\|_{\text{Fro}}$, and $\|A\|_{\text{Spec}}$ correspond to its maximal eigenvalue, Frobenius norm, and spectral norm, respectively. Given a set V, we denote its closure by $\text{cl}(V)$. We abuse notation slightly by denoting $\nabla f(x)$ for both the gradient of function f at x if f is differentiable and a subgradient of f at x, even if f is not differentiable. If f is of the form $f(x, u)$, then $\nabla_x f(x, u)$ denotes the subgradient of f at x while keeping the other variables fixed at u.

2.1 Robust Feasibility Problem

Consider a convex *deterministic* or *nominal* mathematical program

$$\min_x \left\{ f^0(x) : \ x \in X, \ \ f^i(x, u^i) \le 0, \ \forall i \in [m] \right\}, \tag{1}$$

where the domain $X \subset \mathbb{R}^n$ is closed and convex, the functions $f^0(x)$ and $f^i(x, u^i)$ for $i \in [m]$ are convex functions of x, and $u = (u^1, \ldots, u^m)$ is a fixed parameter vector. Without loss of generality we assume the objective function $f^0(x)$ does not have uncertainty. The *robust convex optimization*

problem associated with (1) is

$$\text{Opt} := \min_x \left\{ f^0(x) : \ x \in X, \quad \sup_{u^i \in U^i} f^i(x, u^i) \leq 0, \ \forall i \in [m] \right\}, \tag{2}$$

where $U^1, \ldots, U^m$ are the *uncertainty sets* given for the parameter u^i of constraint $i \in [m]$. Because we assume formulation (1) is convex, the overall optimization problem in (2) is convex.

In this paper, we work under the following mild regularity assumption:

Assumption 2.1. *The constraint functions $f^i(x, u^i)$ for all $i \in [m]$ are finite-valued on the domain $X \times U^i$, convex in x and concave in u^i; X, the domain for x, is closed and convex; and U^i, the domains for u^i, are closed and bounded.*

We take Assumption 2.1 as given for all our results and proofs. Without loss of generality, we assume that the uncertainty set has a Cartesian product form $U^1 \times \ldots \times U^m$, see e.g., [8]; we let $U = U^1 \times \ldots \times U^m$ and write $u = [u^1; \ldots; u^m] \in U$. We do not further assume that the sets U^i are convex. However, for some FOMs we consider, convexity of U^i for $i \in [m]$ will be required.

A convex optimization problem can be solved by solving a polynomial number of associated feasibility problems in a standard way, via a binary search over its optimal value. In particular, let $[\underline{v}_0, \overline{v}_0]$ be an initial interval containing the optimal value of (2). At each iteration k of the binary search, we update the domain $X_k := X \cap \{x : \ f^0(x) \leq v_k\}$ for some $v_k \in [\underline{v}_k, \overline{v}_k]$ and arrive at the following robust feasibility problem:

$$\text{find} \quad x \in X_k \quad \text{s.t.} \quad \sup_{u^i \in U^i} f^i(x, u^i) \leq 0 \quad \forall i \in [m]. \tag{3}$$

Then based on the feasibility/infeasibility status of (3), we update our range $[\underline{v}_{k+1}, \overline{v}_{k+1}]$ and go to iteration $k+1$. In this scheme, we are guaranteed to find a solution $x^* \in X$ whose objective value is within $\delta > 0$ of the optimum value of (2) within at most $\left\lfloor \log_2 \left(\frac{\overline{v}_0 - \underline{v}_0}{\delta} \right) \right\rfloor$ iterations. Therefore, one can equivalently study the complexity of solving robust feasibility problem (3) as opposed to (2). Hence, we focus on solving robust feasibility problem and assume that the constraint on the objective function $f^0(x)$ is already included in the domain X for simplicity in our notation.

Given functional constraints $f^i(x) \leq 0$, $i \in [m]$, most convex optimization methods will declare infeasibility or return an approximate solution $x \in X$ such that $f^i(x) \leq \epsilon$ for $i \in [m]$ for some tolerance level $\epsilon > 0$. Therefore, we will consider the following *robust approximate feasibility problem*:

$$\begin{cases} \textit{Either}\text{: find} \quad x \in X \quad \text{s.t.} \quad \sup_{u^i \in U^i} f^i(x, u^i) \leq \epsilon \quad \forall i \in [m]; \\ \textit{or}\text{: declare infeasibility, } \forall x \in X, \ \exists i \in [m] \quad \text{s.t.} \quad \sup_{u^i \in U^i} f^i(x, u^i) > 0. \end{cases} \tag{4}$$

Then we refer to any feasible solution x to (4), i.e., $x \in X$ such that $\sup_{u^i \in U^i} f^i(x, u^i) \leq \epsilon$ holds for all $i \in [m]$ as a *robust ϵ-feasibility certificate*. Similarly, any realization of the uncertain parameters $\bar{u} \in U$ such that there exists no $x \in X$ satisfying $f^i(x, \bar{u}^i) \leq 0$ for all $i \in [m]$ is referred to as a *robust infeasibility certificate*.

2.2 Saddle Point Problems

Saddle point (SP) problems play a vital role in our developments. In its most general form, a convex-concave SP problem is given by

$$\mathrm{SV} = \inf_{x \in X} \sup_{y \in Y} \phi(x, y) \tag{$\mathcal{S}$}$$

where the function $\phi(x, y)$ is convex in x and concave in y and the domains X, Y are nonempty closed convex sets in Euclidean spaces $\mathbb{E}_x, \mathbb{E}_y$.

Any convex-concave SP problem ($\mathcal{S}$) gives rise to two convex optimization problems that are dual to each other:

$$\begin{array}{rcll} \mathrm{Opt}(P) & = & \inf_{x \in X}[\overline{\phi}(x) := \sup_{y \in Y} \phi(x, y)] & (P) \\ \mathrm{Opt}(D) & = & \sup_{y \in Y}[\underline{\phi}(y) := \inf_{x \in X} \phi(x, y)] & (D) \end{array}$$

with $\mathrm{Opt}(P) = \mathrm{Opt}(D) = \mathrm{SV}$. It is well-known that the solutions to ($\mathcal{S}$) — the saddle points of ϕ on $X \times Y$ — are exactly the pairs $[x; y]$ formed by optimal solutions to the problems (P) and (D).

We quantify the accuracy of a candidate solution $[\bar{x}, \bar{y}]$ to SP problem ($\mathcal{S}$) with the *saddle point gap* given by

$$\epsilon^{\phi}_{\mathrm{sad}}(\bar{x}, \bar{y}) := \overline{\phi}(\bar{x}) - \underline{\phi}(\bar{y}) = \underbrace{\left[\overline{\phi}(\bar{x}) - \mathrm{Opt}(P)\right]}_{\geq 0} + \underbrace{\left[\mathrm{Opt}(D) - \underline{\phi}(\bar{y})\right]}_{\geq 0}. \tag{5}$$

Because convex-concave SP problems are simply convex optimization problems, they can in principle be solved by polynomial-time interior point methods (IPMs). However, the computational complexity of such methods depends heavily on the dimension of the problem. Thus, scalability of resulting algorithms becomes an issue in large-scale applications. As a result, for large-scale SP problems, one has to resort to first-order subgradient-type methods. On a positive note, there are many efficient FOMs for convex-concave SP problems. These in particular include Nesterov's accelerated gradient descent algorithm [32] and Nemirovski's Mirror-Prox algorithm [30], both of which bound the saddle point gap at a rate of $\epsilon^{\phi}_{\mathrm{sad}}(\bar{x}_T, \bar{u}_T) \leq O\left(\frac{1}{T}\right)$ where $\bar{x}_T, \bar{u}_T$ are solutions obtained after T iterations.

2.3 Online Convex Optimization Tools

In order to derive an efficient approach to solving robust feasibility problems, we will employ tools from the online convex optimization domain, which we now briefly outline. We refer to [15, 18, 35] for further details and applications of OCO.

OCO is used to capture decision making in dynamic environments. We are given a finite time horizon T, closed, bounded, and convex domain Z, and in each time period $t \in [T]$, a convex loss function $f_t : Z \to \mathbb{R}$ is revealed. At time periods $t \in [T]$ we must choose a decision $z_t \in Z$, and based on this we suffer a loss of $f_t(z_t)$ and receive some feedback typically in the form of first-order information on f_t. Our goal is to minimize the *weighted regret*

$$\sum_{t=1}^{T} \theta_t f_t(z_t) - \inf_{z \in Z} \sum_{t=1}^{T} \theta_t f_t(z), \tag{6}$$

where $\theta \in \Delta_T$ is a vector of convex combination weights. Note that in the OCO literature, regret is usually defined with uniform weights $\theta_t = 1/T$; recently, weighted regret with general convex

combination weights θ_t was introduced in [20] to allow speedups in convergence rates under certain assumptions. These speedups will be useful for improving the convergence rates of our methods for robust feasibility. Our developments are also based on another important OCO concept of *online saddle point (SP) problem* which was introduced in [20]. In this setting, our domain decomposes as $Z = X \times Y$, and we receive a scoring function $\phi_t : Z \to \mathbb{R}$ where ϕ_t is convex in $x \in X$ and concave in $y \in Y$. At each time period t, our goal is to choose a decision $z_t = [x_t; y_t]$, but instead of the weighted regret, we want to minimize the *weighted online SP gap* defined as

$$\sup_{y \in Y} \sum_{t=1}^{T} \theta_t \phi_t(x_t, y) - \inf_{x \in X} \sum_{t=1}^{T} \theta_t \phi_t(x, y_t). \tag{7}$$

We assume that at each iteration t, the decisions and actions (queries made to the function ϕ_t) of each player, i.e., x_t etc., are revealed to the other and vice versa *immediately after* they make their decision or action. This revealed information from iteration t can then be used by both players in their subsequent decisions and actions in the same period t or in future rounds $t+1$ and so on.

Remark 2.1. A critical feature of OCO is its *non-anticipatory* nature: the current decision z_t must be made only with information from the previous $t-1$ time steps, since the score function f_t (or ϕ_t) is only revealed *after* we choose z_t. For OCO problems arising in the context of RO feasibility problems, it is possible to relax the non-anticipativity assumption of standard OCO and improve overall convergence rates. In particular, in our context, we are allowed to also consider methods to choose z_t in an *anticipatory* manner with knowledge of the scoring function f_t (or ϕ_t); we elaborate more on this in Remark 4.1. ∎

Most OCO algorithms are closely related to offline iterative FOMs. Within the proximal setup of [25], [20] presents online extensions of Mirror Descent and Mirror Prox algorithms to choose the sequence $\{z_t\}_{t=1}^T$ which ensure that the quantities (6) and (7) converge to 0 as $T \to \infty$. These are critical for our developments so we briefly summarize them here. We make the following assumption on Z for the existence of a proximal setup; see [20] for further details.

Assumption 2.2. *Let $\mathbb{E}_z$ be the Euclidean space containing Z. There exists norms $\|\cdot\|, \|\cdot\|_*$ on $\mathbb{E}_z$ that are dual to each other, a distance-generating function $\omega : Z \to \mathbb{R}$ which is 1-strongly convex with respect to $\|\cdot\|$ and leads to an easy-to-compute prox function $\mathrm{Prox}_z(\xi) := \arg\min_{w \in Z} \{\langle \xi, w \rangle + \omega(w) - \langle \omega'(z), w - z \rangle\}$ and set width $\Omega := \max_{z \in Z} \omega(z) - \min_{z \in Z} \omega(z)$ which is finite when Z is bounded.*

The proximal setup of Assumption 2.2 allows us to adjust to the geometry of domain Z. The standard basic domains satisfying Assumption 2.2 include simplex, Euclidean ball, and spectahedron; see [25, Section 1.7] for the standard customizations of proximal setup, i.e., Assumption 2.2, for these basic domains in terms of selection of $\|\cdot\|$ and resulting ω, Prox computation, and set width Ω. In the case of a decomposable domain $Z = X \times Y$, we can build a proximal setup for Z from the individual proximal setups on X and Y. We refer the reader to the references [25, Section 1.7.2] and [26, Section 2.3.3] for further details to handle decomposable domains.

Under Assumption 2.2 and various structural properties, OFO algorithms from [20] achieve the following rates. Theorems 2.1 and 2.2 cover the case of bounding weighted regret and online SP gap respectively under the most basic convexity assumption.

Theorem 2.1 ([20, Theorem 1]). *Suppose there exists $G \in (0,\infty)$ such that $\|\nabla f_t(z)\|_* \le G$ for all $z \in Z$, $t \in [T]$. Let $\theta \in \Delta_T$ be given convex combination weights. Then there exists a* non-anticipatory *OFO algorithm to choose $\{z_t\}_{t=1}^T$ that guarantees*

$$\sum_{t=1}^T \theta_t f_t(z_t) - \inf_{z\in Z} \sum_{t=1}^T \theta_t f_t(z) \le \sqrt{2\Omega \left(\sup_{t\in[T]} \theta_t^2\right) G^2 T}.$$

Theorem 2.2 ([20, Theorem 3]). *Suppose there exists $G \in (0,\infty)$ such that $\|[\nabla_x \phi_t(x,y); -\nabla_y \phi_t(x,y)]\|_* \le G$ for all $[x;y] \in Z$, $t \in [T]$. Let $\theta \in \Delta_T$ be given convex combination weights. Then there exists a* non-anticipatory *OFO algorithm to choose $\{x_t, y_t\}_{t=1}^T$ that guarantees*

$$\sup_{y\in Y} \sum_{t=1}^T \theta_t \phi_t(x_t, y) - \inf_{x\in X} \sum_{t=1}^T \theta_t \phi_t(x, y_t) \le \sqrt{2\Omega \left(\sup_{t\in[T]} \theta_t^2\right) G^2 T}.$$

When the functions f_t are in a sense strongly convex, we obtain the following accelaration.

Theorem 2.3 ([20, Theorem 2]). *Suppose there exists $G \in (0,\infty)$ such that $\|\nabla f_t(z)\|_* \le G$ for all $z \in Z$, $t \in [T]$, and $\alpha > 0$ such that $f_t(z) - \alpha\,\omega(z)$ is convex for all $t \in [T]$. Let $\theta \in \Delta_T$ be the* specific *choice of convex combination weights $\theta_t = 2t/(T^2+T)$. Then there exists a* non-anticipatory *OFO algorithm to choose $\{z_t\}_{t=1}^T$ that guarantees*

$$\sum_{t=1}^T \theta_t f_t(z_t) - \inf_{z\in Z} \sum_{t=1}^T \theta_t f_t(z) \le \frac{2G^2}{\alpha(T+1)}.$$

When the functions f_t are α-strongly convex, i.e., $f_t(x) - \alpha\|x\|_2^2/2$ is convex, we can simply take d.g.f. ω in Theorem 2.3 to be $1/2\|x\|_2^2$; see also [20, Remark 5].

In contrast to Theorems 2.1 and 2.2, note that the convergence rate in Theorem 2.3 depends on the specific choice of weights $\theta_t = 2t/(T^2+T)$.

For the weighted online SP gap, if we assume a smoothness property for the functions ϕ_t, we obtain the following speedup.

Theorem 2.4 ([20, Theorem 5]). *Suppose there exists $L \in (0,\infty)$ such that for all $[x;y], [x';y'] \in Z$, $t \in [T]$, we have*

$$\|[\nabla_x \phi_t(x,y); -\nabla_y \phi_t(x,y)] - [\nabla_x \phi_t(x',y'); -\nabla_y \phi_t(x',y')]\|_* \le L\|[x;y] - [x';y']\|.$$

Let $\theta \in \Delta_T$ be given convex combination weights. Then there exists an anticipatory *OFO algorithm to choose $\{x_t, y_t\}_{t=1}^T$ that guarantees*

$$\sup_{y\in Y} \sum_{t=1}^T \theta_t \phi_t(x_t, y) - \inf_{x\in X} \sum_{t=1}^T \theta_t \phi_t(x, y_t) \le \Omega L \sup_{t\in T} \theta_t.$$

Remark 2.2. For Theorems 2.1, 2.2, and 2.4, the optimal choice of convex combination weights is uniform $\theta_t = 1/T$, which results in overall bounds of $O(1/\sqrt{T})$, $O(1/\sqrt{T})$, and $O(1/T)$ respectively. The only convergence rate that depends on a different selection is Theorem 2.3. However, even if we use the specific selection $\theta_t = 2t/(T^2+T)$, the convergence rates for Theorems 2.1, 2.2, and 2.4 only differ from the optimal by factor of 2. This point on the selection of weights θ_t becomes crucial since our proposed methods in Section 4 will combine two OFO algorithms which must have the *same selection* of weights θ. See Section 4.1 and Table 1 for details. ∎

3 General Framework for Robust Feasibility Problems

In this section, we build a general framework to solve the robust feasibility problem (4) by working with its natural saddle point formulation.

Given constraint functions $f^i(x, u^i)$, $i \in [m]$, let us define $\Phi(x, u) := \max_{i\in[m]} f^i(x, u^i)$. Then $\Phi(x, u)$ is a convex function of x, but not necessarily concave in u. In addition, with this definition of $\Phi(\cdot)$, the robust approximate feasibility problem (4) is equivalent to simply verifying

$$\inf_{x\in X} \sup_{u\in U} \Phi(x, u) = \inf_{x\in X} \max_{i\in[m]} \sup_{u^i\in U^i} f^i(x, u^i) \le \epsilon, \quad \text{or} \quad \inf_{x\in X} \sup_{u\in U} \Phi(x, u) > 0, \tag{8}$$

which is nothing but solving a specific SP problem and checking its value. Analogous to the convex-concave SP gap (5), for a given solution $[\bar{x}, \bar{u}]$, we define the SP gap of problem (8) as

$$\epsilon^{\Phi}_{\text{sad}}(\bar{x}, \bar{u}) := \overline{\Phi}(\bar{x}) - \underline{\Phi}(\bar{u}) = \sup_{u\in U} \Phi(\bar{x}, u) - \inf_{x\in X} \Phi(x, \bar{u}).$$

In general, solving a convex-nonconcave SP problem of form (8), i.e., finding a solution $[\bar{x}, \bar{u}]$ such that $\epsilon^{\Phi}_{\text{sad}}(\bar{x}, \bar{u}) \le \epsilon$, can be difficult. That said, a bound on the SP gap $\epsilon^{\Phi}_{\text{sad}}(\bar{x}, \bar{u})$ along with the value of $\Phi(\bar{x}, \bar{u})$ leads to robust feasibility certificates for (8) as follows.

Theorem 3.1. *Let $\Psi : X \times U \to \mathbb{R}$ be a given function associated with an SP (not necessarily admitting a convex-concave structure). Suppose we have $\bar{x} \in X$, $\bar{u} \in U$, and $\tau \in (0, 1)$ such that $\epsilon^{\Psi}_{\text{sad}}(\bar{x}, \bar{u}) \le \tau\epsilon$. Then if $\Psi(\bar{x}, \bar{u}) \le (1 - \tau)\epsilon$, we have $\sup_{u\in U} \Psi(\bar{x}, u) \le \epsilon$. Moreover, if $\Psi(\bar{x}, \bar{u}) > (1 - \tau)\epsilon$ and $\tau \le \frac{1}{2}$, we have $\inf_{x\in X} \Psi(x, \bar{u}) > 0$.*

Proof. Suppose $\Psi(\bar{x}, \bar{v}) \le (1 - \tau)\epsilon$. Because $\epsilon^{\Psi}_{\text{sad}}(\bar{x}, \bar{v}) = \sup_{u\in U} \Psi(\bar{x}, u) - \inf_{x\in X} \Psi(x, \bar{u}) \le \tau\epsilon$, we have $\sup_{u\in U} \Psi(\bar{x}, u) \le \inf_{x\in X} \Psi(x, \bar{u}) + \tau\epsilon \le \Psi(\bar{x}, \bar{u}) + \tau\epsilon \le \epsilon$. On the other hand, when $\Psi(\bar{x}, \bar{u}) > (1 - \tau)\epsilon$, we have $(1 - \tau)\epsilon < \Psi(\bar{x}, \bar{u}) \le \sup_{u\in U} \Psi(\bar{x}, u) \le \inf_{x\in X} \Psi(x, \bar{u}) + \tau\epsilon$, which implies $\inf_{x\in X} \sup_{u\in U} \Psi(x, u) \ge \inf_{x\in X} \Psi(x, \bar{u}) > (1 - 2\tau)\epsilon \ge 0$ when $\tau \le \frac{1}{2}$. □

Remark 3.1. When $m = 1$, $\Phi(x, u) = f^1(x, u^1)$, and it is thus convex in x and concave in u due to Assumption 2.1. Therefore, in the case of a single robust constraint, i.e., $m = 1$, under Assumption 2.1 and assuming $U = U^1$ is a closed convex set, the optimization problem in (8) reduces to a standard convex-concave SP problem. ∎

While it is not very common, a few robust convex optimization problems come with a single robust constraint and convex uncertainty set U; see for example [2] for a robust version of a support vector machine problem with one constraint. In such cases, based on Remark 3.1, the resulting convex-concave SP problems can directly be solved via efficient FOMs. On the other hand, in the presence of multiple constraints, the function $\Phi(x, u)$ is not concave in $u = [u^1; \dots; u^m]$ even under Assumption 2.1. Nevertheless, when $m > 1$, it is still possible to have a convex-concave SP reformulation of the optimization problem in (8) in an extended space via perspective transformations; we present this in Appendix A. While this reformulation has the benefit of reducing the robust feasibility problem to a well-known and well-studied problem, it destroys the simplicity of the original domains and constraint functions and hence comes with some challenges. Therefore, we develop a framework where we work directly with the convex-nonconcave SP formulation in (8) in the space of original variables. Moreover, because we work in the original space of variables, we simply utilize the first-order information on the original constraint functions f^i and original domains X and U^i. This direct approach in particular allows us to take greater advantage of the structure

of the original formulation such as the availability of efficient projection (prox) computations over domains X, U^i, and/or better parameters for smoothness, Lipschitz continuity, etc., of the functions f^i.

The nonconcavity of $\Phi(x,u)$ in u prevents us from directly applying traditional FOMs for solving convex-concave SP problems to bound the SP gap $\epsilon^{\Phi}_{\text{sad}}(\bar{x},\bar{u})$. However, we show that by just *partially* upper bounding $\epsilon^{\Phi}_{\text{sad}}(\bar{x},\bar{u})$, we can derive a general iterative framework to obtain robust feasibility/infeasibility certificates. In Section 4, we will explore how to apply this framework in practice. Henceforth we will no longer use the shorthand notation $\Phi(x,u) = \max_{i\in[m]} f^i(x,u^i)$, but we will denote the SP gap $\epsilon^{\Phi}_{\text{sad}}(\bar{x},\bar{u})$ as

$$\epsilon(\bar{x},\bar{u}) := \epsilon^{\Phi}_{\text{sad}}(\bar{x},\bar{u}) = \max_{i\in[m]} \sup_{u^i\in U^i} f^i(\bar{x},u^i) - \inf_{x\in X} \max_{i\in[m]} f^i(x,\bar{u}^i). \tag{9}$$

The robust feasibility certificate result from Theorem 3.1 indicates the importance of bounding the SP gap $\epsilon(\bar{x},\bar{u})$. Often, FOMs achieve this by iteratively generating points $x_t \in X$, $u_t \in U$ for $t \in [T]$ and tracking the points $\bar{x}$ and $\bar{u}$ obtained from a convex combination of $\{x_t,u_t\}_{t=1}^T$. In order to simplify our notation, given convex combination weights $\theta \in \Delta_T$ and points $\{x_t,u_t\}_{t=1}^T$, we let

$$\bar{x}_T := \sum_{t=1}^T \theta_t x_t \quad \text{and} \quad \bar{u}_T := \sum_{t=1}^T \theta_t u_t.$$

We now present an upper bound on $\epsilon(\bar{x}_T,\bar{u}_T)$ that follows naturally from the convex-concave structure of functions f^i. To this end, given a set of vectors $y_t \in \Delta_m$ for $t \in [T]$, we also define

$$\epsilon^{\circ}(\{x_t,u_t,\theta_t\}_{t=1}^T) := \max_{i\in[m]} \left\{ \sup_{u^i\in U^i} \sum_{t=1}^T \theta_t f^i(x_t,u^i) - \sum_{t=1}^T \theta_t f^i(x_t,u_t^i) \right\}, \quad \text{and}$$

$$\epsilon^{\bullet}(\{x_t,u_t,y_t,\theta_t\}_{t=1}^T) := \max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t,u_t^i) - \inf_{x\in X} \sum_{t=1}^T \theta_t \sum_{i=1}^m y_t^{(i)} f^i(x,u_t^i),$$

together with

$$\widehat{\epsilon}(\{x_t,u_t,y_t,\theta_t\}_{t=1}^T) := \inf_{x\in X} \sum_{t=1}^T \theta_t \sum_{i=1}^m y_t^{(i)} f^i(x,u_t^i) - \inf_{x\in X} \max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x,u_t^i).$$

Our next result relates these quantities to the value of the SP gap function $\epsilon\left(\bar{x}_T,\bar{u}_T\right)$.

Proposition 3.1. *Let $x_t \in X$ and $u_t \in U$ for $t \in [T]$ be given a set of vectors. Then for any set of vectors $y_t \in \Delta_m$ for $t \in [T]$ and any $\theta \in \Delta_T$, we have*

$$\epsilon\left(\sum_{t=1}^T \theta_t x_t, \sum_{t=1}^T \theta_t u_t\right) \le \epsilon^{\circ}(\{x_t,u_t,\theta_t\}_{t=1}^T) + \epsilon^{\bullet}(\{x_t,u_t,y_t,\theta_t\}_{t=1}^T) + \widehat{\epsilon}(\{x_t,u_t,y_t,\theta_t\}_{t=1}^T). \tag{10}$$

Proof. Given $y_t \in \Delta_m$ for $t \in [T]$ and $\theta \in \Delta_T$, let us define $\bar{x} := \sum_{t=1}^T \theta_t x_t$ and $\bar{u} := \sum_{t=1}^T \theta_t u_t$. We first partition $\epsilon(\bar{x},\bar{u})$ as $\epsilon(\bar{x},\bar{u}) = \overline{\epsilon}(\bar{x},\bar{u}) + \underline{\epsilon}(\bar{x},\bar{u})$ where

$$\overline{\epsilon}(\bar{x},\bar{u}) := \max_{i\in[m]} \sup_{u^i\in U^i} f^i(\bar{x},u^i) - \inf_{x\in X} \max_{i\in[m]} \sup_{u^i\in U^i} f^i(x,u^i),$$

$$\underline{\epsilon}(\bar{x},\bar{u}) := \inf_{x\in X} \max_{i\in[m]} \sup_{u^i\in U^i} f^i(x,u^i) - \inf_{x\in X} \max_{i\in[m]} f^i(x,\bar{u}^i),$$

and then derive upper bounds on $\overline{\epsilon}(\bar{x},\bar{u})$ and $\underline{\epsilon}(\bar{x},\bar{u})$.

We start with bounding $\overline{\epsilon}(\bar{x},\bar{u})$. Because the functions $f^i(x,u^i)$ are convex in x for all i and $\theta\in\Delta_T$, we have $\max_{i\in[m]}\sup_{u^i\in U^i} f^i(\bar{x},u^i)\le \max_{i\in[m]}\sup_{u^i\in U^i}\sum_{t=1}^T\theta_t f^i(x_t,u^i)$. Therefore,

$$\begin{aligned}
\overline{\epsilon}(\bar{x},\bar{u}) &= \max_{i\in[m]}\sup_{u^i\in U^i} f^i(\bar{x},u^i) - \inf_{x\in X}\max_{i\in[m]}\sup_{u^i\in U^i} f^i(x,u^i)\\
&\le \max_{i\in[m]}\sup_{u^i\in U^i}\sum_{t=1}^T\theta_t f^i(x_t,u^i) - \max_{i\in[m]}\sum_{t=1}^T\theta_t f^i(x_t,u_t^i) + \max_{i\in[m]}\sum_{t=1}^T\theta_t f^i(x_t,u_t^i) - \inf_{x\in X}\max_{i\in[m]}\sup_{u^i\in U^i} f^i(x,u^i)\\
&\le \max_{i\in[m]}\left\{\sup_{u^i\in U^i}\sum_{t=1}^T\theta_t f^i(x_t,u^i) - \sum_{t=1}^T\theta_t f^i(x_t,u_t^i)\right\} + \max_{i\in[m]}\sum_{t=1}^T\theta_t f^i(x_t,u_t^i) - \inf_{x\in X}\max_{i\in[m]}\sup_{u^i\in U^i} f^i(x,u^i),
\end{aligned}\tag{11}$$

where the last inequality follows since $\max_{i\in[m]}\{\alpha_i-\beta_i\}\ge\max_{i\in[m]}\alpha_i-\max_{i\in[m]}\beta_i$ for any sequence of numbers α_i,β_i, $i\in[m]$.

Note that $\inf_{x\in X}\max_{i\in[m]} f^i(x,u^i)\ge\inf_{x\in X}\max_{i\in[m]}\sum_{t=1}^T\theta_t f^i(x,u^i)$ because under Assumption 2.1 the functions $f^i(x,u^i)$ are concave in u^i for all i. Thus, we arrive at

$$\begin{aligned}
\underline{\epsilon}(\bar{x},\bar{u}) &= \inf_{x\in X}\max_{i\in[m]}\sup_{u^i\in U^i} f^i(x,u^i) - \inf_{x\in X}\max_{i\in[m]} f^i(x,\bar{u}^i)\\
&\le \inf_{x\in X}\max_{i\in[m]}\sup_{u^i\in U^i} f^i(x,u^i) - \inf_{x\in X}\sum_{t=1}^T\theta_t\sum_{i=1}^m y_t^{(i)} f^i(x,u_t^i) + \inf_{x\in X}\sum_{t=1}^T\theta_t\sum_{i=1}^m y_t^{(i)} f^i(x,u_t^i)\\
&\qquad - \inf_{x\in X}\max_{i\in[m]}\sum_{t=1}^T\theta_t f^i(x,u_t^i)\\
&= \inf_{x\in X}\max_{i\in[m]}\sup_{u^i\in U^i} f^i(x,u^i) - \inf_{x\in X}\sum_{t=1}^T\theta_t\sum_{i=1}^m y_t^{(i)} f^i(x,u_t^i) + \widehat{\epsilon}(\{x_t,u_t,y_t,\theta_t\}_{t=1}^T).
\end{aligned}\tag{12}$$

Then by summing (11) and (12) and rearranging the terms, we deduce the result. □

We are now ready to state our main result, which is analogous to Theorem 3.1 except that we do not need to bound all three terms in (10), but instead it suffices to guarantee that

$$\epsilon^\circ(\{x_t,u_t,\theta_t\}_{t=1}^T) + \epsilon^\bullet(\{x_t,u_t,y_t,\theta_t\}_{t=1}^T)\le\epsilon.$$

When this holds, based on the value of $\max_{i\in[m]}\sum_{t=1}^T\theta_t f^i(x_t,u_t^i)$ we can then obtain robust ϵ-feasibility/infeasibility certificates.

Theorem 3.2. *Suppose we have sequences $\{x_t,u_t,y_t,\theta_t\}_{t=1}^T$ with $x_t\in X$, $u_t\in U$, $y_t\in\Delta_m$ for all $t\in[T]$, $\theta\in\Delta_T$. Let $\tau\in(0,1)$. If $\epsilon^\circ(\{x_t,u_t,\theta_t\}_{t=1}^T)\le\tau\epsilon$ and $\max_{i\in[m]}\sum_{t=1}^T\theta_t f^i(x_t,u_t^i)\le(1-\tau)\epsilon$, then the solution $\bar{x}_T:=\sum_{t=1}^T\theta_t x_t$ is ϵ-feasible with respect to (4). If $\epsilon^\bullet(\{x_t,u_t,y_t,\theta_t\}_{t=1}^T)\le(1-\tau)\epsilon$ and $\max_{i\in[m]}\sum_{t=1}^T\theta_t f^i(x_t,u_t^i)>(1-\tau)\epsilon$, then (4) is infeasible.*

Proof. First suppose there exists a $\tau \in (0,1)$ and corresponding vectors $\{x_t, u_t, y_t, \theta_t\}_{t=1}^T$ such that $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) \le \tau\epsilon$ and $\max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \le (1-\tau)\epsilon$ holds as well. Note that

$$\begin{aligned}\tau\epsilon \ge \epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) &= \max_{i\in[m]} \left\{ \sup_{u^i\in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \right\} \\ &\ge \max_{i\in[m]} \sup_{u^i\in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i), \end{aligned} \tag{13}$$

where the last inequality follows since $\max_{i\in[m]}\{\alpha_i - \beta_i\} \ge \max_{i\in[m]} \alpha_i - \max_{i\in[m]} \beta_i$ for any sequence of numbers α_i, β_i, $i \in [m]$. Then $\bar{x}_T$ is an ϵ-feasible solution for (4) because

$$\begin{aligned}\max_{i\in[m]} \sup_{u^i\in U^i} f^i\left(\bar{x}_T, u^i\right) &= \max_{i\in[m]} \sup_{u^i\in U^i} f^i\left(\sum_{t=1}^T \theta_t x_t, u^i\right) \\ &\le \max_{i\in[m]} \sup_{u^i\in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) \le \tau\epsilon + \max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \le \epsilon,\end{aligned}$$

where the first inequality follows from the convexity of the functions f^i and the fact that $\theta \in \Delta_T$, the second inequality from (13), and the last inequality holds since $\max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \le (1-\tau)\epsilon$.

On the other hand, suppose $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) \le (1-\tau)\epsilon$ and $\max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) > (1-\tau)\epsilon$. Note that

$$\begin{aligned}\inf_{x\in X} \sum_{t=1}^T \theta_t \sum_{i=1}^m y_t^{(i)} f^i(x, u_t^i) &\le \inf_{x\in X} \sum_{t=1}^T \theta_t \max_{i\in[m]} f^i(x, u_t^i) \\ &\le \inf_{x\in X} \sum_{t=1}^T \theta_t \max_{i\in[m]} \sup_{u^i\in U^i} f^i(x, u^i) = \inf_{x\in X} \max_{i\in[m]} \sup_{u^i\in U^i} f^i(x, u^i),\end{aligned} \tag{14}$$

where the first inequality follows since $y_t \in \Delta_m$ for all $t \in [T]$, the second inequality holds because $f^i(x, u_t^i) \le \sup_{u^i\in U^i} f^i(x, u^i)$ for all $i \in [m]$ and $y_t^{(i)} \ge 0$ for $i \in [m]$, $t \in [T]$, and the last equation follows from $\theta \in \Delta_T$. Then using the bound

$$(1-\tau)\epsilon \ge \epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) = \max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) - \inf_{x\in X} \sum_{t=1}^T \theta_t \sum_{i=1}^m y_t^{(i)} f^i(x, u_t^i), \tag{15}$$

we arrive at

$$\inf_{x\in X} \max_{i\in[m]} \sup_{u^i\in U^i} f^i(x, u^i) \ge \inf_{x\in X} \sum_{t=1}^T \theta_t \sum_{i=1}^m y_t^{(i)} f^i(x, u_t^i) \ge \max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) - (1-\tau)\epsilon > 0,$$

where the first inequality follows from inequality (14), the second inequality from (15) and the last inequality holds because $\max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) > (1-\tau)\epsilon$. This implies (4) is infeasible. □

Corollary 3.1. *Suppose $\{x_t, u_t, y_t, \theta_t\}_{t=1}^T$ with $x_t \in X$, $u_t \in U$, $y_t \in \Delta_m$ for all $t \in [T]$, and $\theta \in \Delta_T$ is such that there exists $\kappa^\circ, \kappa^\bullet \in (0,1)$ satisfying $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) \le \epsilon\,\kappa^\circ$ and $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) \le \epsilon\,\kappa^\bullet$ with $\kappa^\circ + \kappa^\bullet \le 1$. Let $\tau \in [\kappa^\circ, 1-\kappa^\bullet]$. Whenever $\max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \le (1-\tau)\epsilon$ as well, the solution $\bar{x}_T := \sum_{t=1}^T \theta_t x_t$ is ϵ-feasible with respect to (4). Also, whenever $\max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) > (1-\tau)\epsilon$, then (4) is infeasible.*

Proof. Note that $\tau \in (0,1)$ follows from its definition, $\kappa^\circ, \kappa^\bullet \geq 0$, and $\kappa^\circ + \kappa^\bullet \leq 1$. Furthermore, the interval $[\kappa^\circ, 1-\kappa^\bullet]$ is well-defined since $\kappa^\circ \leq 1-\kappa^\bullet$ always holds. Moreover, $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \epsilon\,\kappa^\circ \leq \epsilon\,\tau$ and $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) \leq \epsilon\,\kappa^\bullet \leq \epsilon(1-\tau)$ holds from the definition of τ. The result now follows from Theorem 3.2. □

In Remark 3.2 we briefly comment on the additional term $\widehat{\epsilon}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T)$ in (10).

Remark 3.2. Consider the term

$$\widehat{\epsilon}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) = \inf_{x\in X} \sum_{t=1}^T \theta_t \sum_{i=1}^m y_t^{(i)} f^i(x, u_t^i) - \inf_{x\in X} \max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x, u_t^i).$$

When $m = 1$, we have $\widehat{\epsilon}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) = 0$ because $y = y^{(1)} = 1$. In general, when $m \geq 2$, $\widehat{\epsilon}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T)$ cannot be simplified further even under Assumption 2.1 because of the nonconcavity of the problem. ■

Theorem 3.2 and Corollary 3.1 points to our general iterative framework for finding robust feasibility/infeasibility certificates of (4): generate sequences $\{x_t, u_t, y_t\}_{t=1}^T$ iteratively to bound $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$ and $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T)$, and then evaluate the term $\max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i)$. In the next section, we describe some approaches to generate these sequences in practice.

4 Customizations of the General Framework

In this section, we examine how to generate the sequences $\{x_t, u_t, y_t\}_{t=1}^T$ in practice. In Section 4.1, we first interpret the terms $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$ and $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T)$ from Section 3 as weighted regret and weighted online SP gap terms respectively; this gives rise to our online first-order (OFO)-based approach. In Section 4.2, we modify the pessimization oracle-based approach of [29] to solving (4) within our framework. In Section 4.3, we examine the nominal feasibility oracle-based approach of [5] within the context of our general framework. Finally, in Section 4.4, we summarize and compare the convergence rates achievable via various customizations of these different approaches.

4.1 The OFO-based Approach

Let us first consider $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$. For any $i \in [m]$, given x_t, we define the function $f_t^i : U^i \to \mathbb{R}$ as $f_t^i(u^i) = -f^i(x_t, u^i)$. Then the function $f_t^i(u^i)$ is convex in u^i under Assumption 2.1, and the subterm of $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$ given by

$$\sup_{u^i\in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \tag{16}$$

is the weighted regret (6) corresponding to the sequence of functions $\{f_t^i\}_{t=1}^T$. When the uncertainty sets U^i, $i \in [m]$ admit proximal setups as in Assumption 2.2, Theorems 2.1 and 2.3 from Section 2.3 point to efficient online first-order algorithms for choosing $\{u_t^i\}_{i=1}^m$ to bound subterms (16) as weighted regret. Therefore, by choosing T sufficiently large and employing one of these algorithms for each constraint $i \in [m]$, we can generate a sequence $\{u_t\}_{t=1}^T$ that guarantees $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \tau\epsilon$ for *any* sequence $\{x_t\}_{t=1}^T$.

On the other hand, given $u_t^i \in U^i$ for $i \in [m]$, let us define $\phi_t(x, y) := \sum_{i=1}^m y^{(i)} f^i(x, u_t^i)$ for any $y \in \Delta_m$. Then for any fixed $y \in \Delta_m$, $\phi_t(x, y)$ is convex in x over X since f^i are convex in x by

Assumption 2.1, and $\phi_t(x,y)$ is concave (in fact just linear) in y for any fixed $x \in X$. We can then rewrite $\epsilon^\bullet(\{x_t,u_t,y_t,\theta_t\}_{t=1}^T)$ as

$$\begin{aligned}
\epsilon^\bullet(\{x_t,u_t,y_t,\theta_t\}_{t=1}^T) &= \max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t,u_t^i) - \inf_{x\in X} \sum_{t=1}^T \theta_t \sum_{i=1}^m y_t^{(i)} f^i(x,u_t^i) \\
&= \sup_{y\in\Delta_m} \sum_{i=1}^m y^{(i)} \sum_{t=1}^T \theta_t f^i(x_t,u_t^i) - \inf_{x\in X} \sum_{t=1}^T \theta_t \sum_{i=1}^m y_t^{(i)} f^i(x,u_t^i) \\
&= \sup_{y\in\Delta_m} \sum_{t=1}^T \theta_t \sum_{i=1}^m y^{(i)} f^i(x_t,u_t^i) - \inf_{x\in X} \sum_{t=1}^T \theta_t \sum_{i=1}^m y_t^{(i)} f^i(x,u_t^i) \\
&= \sup_{y\in\Delta_m} \sum_{t=1}^T \theta_t \phi_t(x_t,y) - \inf_{x\in X} \sum_{t=1}^T \theta_t \phi_t(x,y_t). \qquad (17)
\end{aligned}$$

This is exactly the weighted online SP gap (7) with functions $\phi_t(x,y)$ and domains X and $Y = \Delta_m$. Therefore, bounding $\epsilon^\bullet(\{x_t,u_t,y_t,\theta_t\}_{t=1}^T)$ exhibits a natural online SP interpretation. Theorems 2.2 and 2.4 from Section 2.3 point to efficient OFO algorithms that choose $[x_t;y_t]$ to bound the online SP gap term (17) when the joint domain $X \times Y$ admits a proximal setup, e.g., Assumption 2.2. Because Y is a simplex domain and hence has a proximal setup, $X \times Y$ admits a proximal setup as long as X admits one. Thus, by choosing T sufficiently large and employing one of these algorithms, we can generate a sequence $\{x_t,y_t\}_{t=1}^T$ to guarantee $\epsilon^\bullet(\{x_t,u_t,y_t,\theta_t\}_{t=1}^T) \le (1-\tau)\epsilon$ for *any* sequence $\{u_t\}_{t=1}^T$.

Our OFO-based approach can thus be described as follows: choose T sufficiently large, then *simultaneously* generate $\{u_t\}_{t=1}^T$ and $\{x_t,y_t\}_{t=1}^T$ iteratively using tools from Section 2.3 to bound $\epsilon^\circ(\{x_t,u_t,\theta_t\}_{t=1}^T) \le \tau\epsilon$ and $\epsilon^\bullet(\{x_t,u_t,y_t,\theta_t\}_{t=1}^T) \le (1-\tau)\epsilon$. Thereby, via Theorem 3.2, we can solve the robust feasibility problem (4) using only cheap OFO algorithms, which avoids relying on a nominal feasibility oracle for x such as the one used in [5].

Remark 4.1. Recall that Theorems 2.1, 2.3, and 2.2 generate the current decision z_t in a *non-anticipatory* manner with only information from the previous $t-1$ iterations. On the other hand, Theorem 2.4 generates $[x_t;y_t]$ in an *anticipatory* manner with information from the current function $\phi_t(x,y)$. In the context of our general framework, it is permissible to use *at most one* anticipatory algorithm to generate either $\{x_t,y_t\}_{t=1}^T$ or $\{u_t\}_{t=1}^T$. This is because if generating $[x_t;y_t]$ requires knowledge of $\phi_t(x,y) = \sum_{i=1}^m y^{(i)} f^i(x,u_t^i)$, then we need to already have generated the points $\{u_t^i\}_{i=1}^m$ to build the function ϕ_t. Of course, we could reverse this, and generate $\{u_t^i\}_{i=1}^m$ in an anticipatory manner, as long as we then use a non-anticipatory algorithm to generate $[x_t;y_t]$. However, it is not possible for *both* $[x_t;y_t]$ and $\{u_t^i\}_{i=1}^m$ to be generated in an anticipatory manner; at least one must be generated by a non-anticipatory algorithm. ∎

Our OFO-based approach is quite flexible in terms of the selection of OFO algorithms, and is certainly not restricted to only using those mentioned in Section 2.3, see Remark 4.8. To demonstrate this, we describe our OFO-based approach for obtaining robust feasibility/infeasibility certificates for (4) precisely in Algorithm 1 for general choices of OFO algorithms. We define some notation. Let $\mathcal{A}_i$ be OFO algorithms that work with domains U^i for each $i \in [m]$ and bound the weighted regret terms (16), and let $\mathcal{A}_{x,y}$ be an OFO algorithm that works with the domain $X \times Y$ where $Y = \Delta_m$ to bound the online SP gap (17).

Algorithm 1 OFO-based approximate robust feasibility solver.

input: tolerance level $\epsilon > 0$, sufficiently large $T = T(\epsilon)$ and convex combination weights $\theta_1, \ldots, \theta_T > 0$ to achieve $\epsilon/2$ weighted regret guarantees from $\mathcal{A}_i$ and $\epsilon/2$ weighted online SP gap guarantee from $\mathcal{A}_{x,y}$.
output: either $\bar{x} \in X$ such that $\sup_{u^i \in U^i} f^i(\bar{x}, u^i) \leq \epsilon$ for all $i \in [m]$, or an infeasibility certificate for (4).
initialize $u_0^i \in U^i$ for $i \in [m]$ and $[x_0; y_0] \in X \times Y = X \times \Delta_m$ according to $\mathcal{A}_i, \mathcal{A}_{x,y}$ respectively.
for $t = 1, \ldots, T$ **do**
  for $i = 1, \ldots, m$ **do**
    update u_t^i according to $\mathcal{A}_i$ using loss functions $f_s^i(u) = -f^i(x_s, u)$ and weights θ_s, $s \in [t-1]$ over U^i.
  end for
  update $[x_t; y_t]$ according to $\mathcal{A}_{x,y}$ which solves online SP problem with convex-concave functions $\phi_s(x, y) = \sum_{i=1}^m y_i f^i(x, u_s^i)$ and weights θ_s, $s \in [t]$ over domain $X \times Y$.
  obtain upper bounds $\epsilon\kappa_t^\circ \geq \epsilon^\circ(\{x_s, u_s, \theta_s\}_{s=1}^t)$, $\epsilon\kappa_t^\bullet \geq \epsilon^\bullet(\{x_s, u_s, y_s, \theta_s\}_{s=1}^t)$ from $\mathcal{A}_i$, $\mathcal{A}_{x,y}$ respectively.

  if $\kappa_t^\circ + \kappa_t^\bullet \leq 1$ **then**
    set $\vartheta_t := \max_{i \in [m]} \sum_{s=1}^t \theta_s f^i(x_s, u_s^i)$.
    set $\tau_t := 1 - \kappa_t^\bullet$.
    if $\vartheta_t > (1 - \tau_t)\epsilon$ **then return** 'infeasible'.
    if $\vartheta_t \leq (1 - \tau_t)\epsilon$ **then return** $\bar{x}_t = \sum_{s=1}^t \theta_s x_s$ as a robust ϵ-feasible solution to (4).
  end if
end for

Remark 4.2. Note that Algorithm 1 chooses $\tau_t = 1 - \kappa_t^\bullet$, whereas Corollary 3.1 allows us to choose from a range $\tau_t \in [\kappa_t^\circ, 1 - \kappa_t^\bullet]$. This is because it is theoretically possible for (4) to simultaneously be infeasible and robust ϵ-feasible, but in practice we would like to discover infeasibility of (4) rather than an approximately feasible solution. Then the best value for $\tau_t \in [\kappa_t^\circ, 1 - \kappa_t^\bullet]$ in detecting infeasibility of (4) is to set $\tau_t = 1 - \kappa_t^\bullet$. ∎

Remark 4.3. Let us briefly discuss our motivation to allow nonuniform weights θ_t and thus work with weighted regret and weighted online SP gap as opposed to the standard regret and online SP gap. Nonuniform weights are especially important when we would like to select specific customization of algorithms for $\mathcal{A}_i$ and $\mathcal{A}_{x,y}$ to exploit structural properties of the constraint functions f^i and achieve better convergence rates, a prime example being the strong convexity assumption of Theorem 2.3. Moreover, our OFO-based approach to solve the robust feasibility problem (4), i.e., Algorithm 1, requires us to run several OFO algorithms in conjunction with each other while using the *same* convex combination weights θ_t in *all* of the bounds optimized by algorithms $\mathcal{A}_i, \mathcal{A}_{x,y}$. In Remark 2.2, we observed that using the specific selection of weights $\theta_t = 2t/(T^2+T)$ does not impact the bounds in Theorems 2.1, 2.2, and 2.4 in terms of their dependence on T. Hence, using Theorem 2.3 together with any of Theorems 2.1, 2.2, and 2.4 will not have an adverse effect on the overall convergence rate of our approach. ∎

4.2 The Pessimization Oracle-Based Approach

At iteration t, the approach of Mutapcic and Boyd [29] generates solutions $x_t \in X$ by solving an extended nominal problem

$$\min_{x \in X} \left\{ f^0(x) : \ f^i(x, u^i) \leq 0, \ \forall u^i \in \hat{U}_{t-1}^i, \ i \in [m] \right\}, \tag{18}$$

where $\hat{U}^i_{t-1} \subset U^i$ are finite approximate uncertainty sets. The noise realizations $\{u^i_t\}_{i=1}^m$ are then generated by calling the *pessimization oracles* on the current solution x_t. More precisely, given $x_t \in X$, the pessimization oracles solve $\sup_{u^i \in U^i} f^i(x_t, u^i)$ and return

$$u^i_t \in U^i \quad \text{s.t.} \quad f^i(x_t, u^i_t) \geq \sup_{u^i \in U^i} f^i(x_t, u^i) - \tau\epsilon. \tag{19}$$

If for all $i \in [m]$ we have $f^i(x_t, u^i_t) \leq (1-\tau)\epsilon$, then we terminate and declare x_t is a robust ϵ-feasible and optimal solution; otherwise, we append $\hat{U}^i_t = \hat{U}^i_{t-1} \cup \{u^i_t\}$ and re-solve (18) with the new approximate sets $\hat{U}^i_t$. We terminate when $\hat{U}^i_t = \hat{U}^i_{t-1}$ for all $i \in [m]$ with a robust optimal solution x_t. It is shown in [29, Section 5.2] that the number of iterations T needed before termination with a robust ϵ-feasible solution x_T is upper bounded by $(1 + O(1/\epsilon))^n$ where n is the dimension of x.

Suppose now that we are interested in robust feasibility (4). In [29, Section 5.3], a number of variations for generating x_t are suggested by modifying (18). In contrast, we propose the following modification: instead of solving (18), we generate $\{x_t, y_t\}_{t=1}^T$ to bound $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T)$. Then the pessimization oracle-based approach fits within our general framework as a special case.

Theorem 4.1. *Let $\tau \in (0,1)$. Suppose $\{x_t, y_t\}_{t=1}^T$ are generated iteratively in a non-anticipatory manner to guarantee that $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) \leq (1-\tau)\epsilon$ for* any *sequence $\{u_t\}_{t=1}^T$. Suppose u^i_t are generated by pessimization oracles (19) for $i \in [m]$. If there exists $t \in [T]$ such that for all $i \in [m]$ we have $f^i(x_t, u^i_t) \leq (1-\tau)\epsilon$, then x_t is a robust ϵ-feasible solution to (4). If $\max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u^i_t) \leq (1-\tau)\epsilon$, then $\bar{x}_T = \sum_{t=1}^T \theta_t x_t$ is a robust ϵ-feasible solution to (4). If $\max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u^i_t) > (1-\tau)\epsilon$, then we certify that (4) is robust infeasible.*

Proof. It is clear that if there exists $t \in [T]$ such that for all $i \in [m]$ we have $f^i(x_t, u^i_t) \leq (1-\tau)\epsilon$, then x_t is a robust ϵ-feasible solution to (4). Furthermore, the fact that $\max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u^i_t) > (1-\tau)\epsilon$ implies robust infeasibility of (4) follows from our assumption that $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) \leq (1-\tau)\epsilon$ and Theorem 3.2. To show that $\max_{i\in[m]} \sum_{t=1}^T f^i(x_t, u^i_t) \leq (1-\tau)\epsilon$ implies that $\bar{x}_T$ is robust ϵ-feasible, we only need to show that $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \tau\epsilon$. Observe that by our definition of u^i_t in (19), we have $f^i(x_t, u^i_t) \geq \sup_{u^i \in U^i} f^i(x_t, u^i) - \tau\epsilon$, hence the regret terms in $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$ satisfy

$$\begin{aligned}
\sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t f^i(x_t, u^i_t) &\leq \sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t \left(\sup_{u^i \in U^i} f^i(x_t, u^i) - \tau\epsilon \right) \\
&= \sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t \sup_{u^i \in U^i} f^i(x_t, u^i) + \tau\epsilon \\
&\leq \tau\epsilon.
\end{aligned}$$

Then

$$\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) = \max_{i\in[m]} \left\{ \sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t f^i(x_t, u^i_t) \right\} \leq \tau\epsilon,$$

and the result follows from Theorem 3.2. ∎

Theorem 4.1 can only be used to certify robust feasibility/infeasibility, hence to find a robust ϵ-optimal solution, we must perform a binary search and solve at most $O(\log(1/\epsilon))$ instances of

robust feasibility problems. Despite this, in Section 4.4, we discuss how using OFO algorithms to generate $\{x_t, y_t\}_{t=1}^T$ as in Section 4.1 results in much better guarantees than using (18) as proposed by [29], even when taking into account the additional $O(\log(1/\epsilon))$ factor.

Remark 4.4. In the pessimization oracle-based approach, the noises u_t need to be generated with knowledge of x_t, because it is not possible to guarantee $f^i(x_t, u_t^i) \geq \sup_{u^i \in U^i} f^i(x_t, u^i) - \tau\epsilon$ if the vectors u_t^i were chosen with only the knowledge of $x_1, \ldots, x_{t-1}$. In contrast, as stated in Remark 4.1, our OFO-based approach could also generate the u_t in a non-anticipatory manner only with the knowledge of previous iterations. This is useful when we would like to accelerate the convergence rate of $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T)$ by allowing $[x_t; y_t]$ to be generated in an anticipatory manner. ∎

4.3 The Nominal Feasibility Oracle-Based Approach

The nominal feasibility oracle-based approach of Ben-Tal et al. [5] suggest using OFO algorithms to choose a sequence $\{u_t\}_{t=1}^T$ that guarantees $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$ is small, in a non-anticipatory fashion, for *any* sequence $\{x_t\}_{t=1}^T$, which is essentially the same as our OFO-based approach outlined in Section 4.1. The key differentiating point between our OFO-based approach and that of [5] lies in how the sequence $\{x_t\}_{t=1}^T$ is chosen. At step t, [5] utilizes a *nominal feasibility oracle*. That is, given parameters u_t, they call a powerful, and potentially expensive, nominal feasibility oracle that solves the following feasibility problem to ϵ-accuracy

$$\begin{cases} \textit{Either: find} \quad x \in X \quad \text{s.t.} \quad f^i(x, u_t^i) \leq \epsilon/2 \quad \forall i \in [m]; \\ \textit{or: declare infeasibility, } \forall x \in X, \ \exists i \in [m] \quad \text{s.t.} \quad f^i(x, u_t^i) > 0. \end{cases} \tag{20}$$

We denote $x_t \in X$ to be the point returned by this oracle at step t, if it exists. For this approach, the outputs of a nominal feasibility oracle can be used to deduce a result similar to Theorem 3.2, except that we no longer need to evaluate $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T)$, we just need to bound $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$.

Theorem 4.2. *Given weights $\theta \in \Delta_T$, suppose that the sequence $\{u_t\}_{t=1}^T$ is generated in a non-anticipatory manner to guarantee $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \epsilon$ for any sequence $\{x_t\}_{t=1}^T$. Furthermore, suppose that at each step $t \in [T]$, x_t is generated by the nominal feasibility oracle which solves (20). If there exists $t \in [T]$ such that (20) declares infeasibility, then (4) is infeasible. Otherwise, if x_t satisfies $f^i(x_t, u_t^i) \leq \epsilon/2$ for all $t \in [T]$ and $i \in [m]$, we have a robust ϵ-feasibility certificate for (4).*

Proof. If (20) declares infeasibility, then it is obvious that the robust feasibility problem is infeasible. We focus on the latter case. By the premise of the theorem, we have $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \epsilon$. Let us evaluate $\max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i)$. Because $\theta \in \Delta_T$ and from the definition of the nominal feasibility oracle we have $f^i(x_t, u_t^i) \leq \epsilon/2$ for all $t \in [T]$ and $i \in [m]$, we conclude $\max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \leq \epsilon/2$. The conclusion now follows from Theorem 3.2 since the assumptions are met for $\tau = 1/2$. □

Thus, the approach of [5], which works with nominal feasibility oracles, fits within our framework right away. We next make three important remarks.

Remark 4.5. Similar to Remark 4.4, it is not possible to generate the vectors $\{u_t^i\}_{i=1}^m$ in an anticipatory manner with the knowledge of x_t, and $[x_t; y_t]$ generated in a non-anticipatory fashion only with the knowledge of the previous iterations. This is because a critical property required of the vectors x_t is that $f^i(x_t, u_t^i) \leq \epsilon/2$, which is not possible to guarantee if x_t were chosen only with the knowledge of $\{u_1^i\}_{i=1}^m, \ldots, \{u_{t-1}^i\}_{i=1}^m$.

Remark 4.6. Theorem 4.2 states that the nominal feasibility oracle-based approach can solve robust feasibility problems (4). This then recovers [5, Theorems 1,2]. In addition, we next make a nice and practical observation that was overlooked in [5]. We show that slightly adjusting this oracle will let us directly solve the robust *optimization* problem (2), i.e., optimize a convex objective function $f^0(x)$ instead of relying on a binary search over the optimal objective value. Recall that Opt denotes the optimal value of the RO problem (see (2)). Naively, to solve for Opt, we would embed f^0 into the constraint set, and then perform a binary search over the robust feasible set by repeatedly applying the oracle-based approach and Theorem 4.2 to check for robust feasibility. Suppose now that we work with a *nominal optimization oracle*. That is, when given fixed parameters u_t, instead of using a nominal feasibility oracle to solve (20), we have access to a nominal optimization oracle that solves

$$\mathrm{Opt}_t = \inf_x \left\{ f^0(x) : f^i(x, u_t^i) \le 0,\ i \in [m], \quad x \in X \right\}.$$

When solving for Opt_t, most nominal convex optimization solvers will either declare that the constraints are infeasible, or return a point $x_t \in X$ such that $f^i(x_t, u_t^i) \le \epsilon/2$ and $f^0(x_t) \le \mathrm{Opt}_t + \epsilon$. It is clear that $f^0(x_t) \le \mathrm{Opt}_t + \epsilon \le \mathrm{Opt} + \epsilon$. Given such a sequence of points $\{x_t\}_{t=1}^T$, from Theorem 4.2 we deduce that $\bar{x}_T = \sum_{t=1}^T \theta_t x_t$ is a robust ϵ-feasible solution. Then by convexity of f^0, we have

$$f^0(\bar{x}_T) \le \sum_{t=1}^T \theta_t f^0(x_t) \le \sum_{t=1}^T \theta_t (\mathrm{Opt} + \epsilon) = \mathrm{Opt} + \epsilon.$$

Hence, not only do we claim that $\bar{x}_T$ is robust ϵ-feasible, but that it is also ϵ-optimal. Thus, when our oracle can return ϵ-optimal solutions, which most solvers can, we eliminate the need to perform a binary search. ∎

Below we elaborate on the differences between Theorem 4.2 and Theorem 3.2.

Remark 4.7. In contrast to Theorem 3.2, Theorem 4.2 does not need to control the term $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T)$. There are two reasons for this: (i) due to (20), each point x_t satisfies $f^i(x_t, u_t^i) \le \epsilon/2$, hence $\max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \le \epsilon/2$ always holds; therefore, the infeasibility part of Theorem 3.2 never becomes relevant; and (ii) due to the oracle solving (20), infeasibility may be declared at any step $t \in [T]$ in Theorem 4.2. This offers the possibility of stopping early rather than having to wait until all T steps are completed. Thus, the nominal feasibility oracle-based approach trades off using more effort at each iteration t to solve (20) for the ability to terminate early. In contrast, our OFO-based approach opts to keep the per-iteration cost cheap while giving up the ability to terminate early. More formally, let us examine a particular way of solving (20) within a nominal feasibility oracle. Note (20) is equivalent to checking $\mathrm{SV}_t \le \epsilon/2$ or $\mathrm{SV}_t > 0$ where

$$\mathrm{SV}_t := \min_{x\in X} \max_{y\in Y} \sum_{i=1}^m y^{(i)} f^i(x, u_t^i), \quad \text{and} \quad Y = \Delta_m.$$

As in Theorem 3.1, this can be solved by finding $x_t \in X$ and $y_t \in Y$ that bounds the SP gap

$$\max_{y\in Y} \sum_{i=1}^m y^{(i)} f^i(x_t, u_t^i) - \inf_{x\in X} \sum_{i=1}^m y_t^{(i)} f^i(x, u_t^i) \le \epsilon/6. \tag{21}$$

Then, by evaluating $\sum_{i=1}^m y_t^{(i)} f^i(x_t, u_t^i)$ and checking whether it is $\le \epsilon/6$ or not, we determine

Table 1: Summary of convergence rates for bounding $\epsilon^{\circ}(\{x_t, u_t, \theta_t\}_{t=1}^T)$ and $\epsilon^{\bullet}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T)$.

Weights θ_t	Assumptions on f^i	$\mathcal{A}_i$	$r_u(\epsilon)$	Assumptions on f^i	$\mathcal{A}_{x,y}$	$r_x(\epsilon)$
$1/T$	Concave in u	Thm 2.1	$O(1/\epsilon^2)$	Convex in x	Thm 2.2	$O(\log(m)/\epsilon^2)$
$2t/(T^2+T)$	Strongly concave in u	Thm 2.3	$O(1/\epsilon)$	Smooth in x	Thm 2.4	$O(\log(m)/\epsilon)$

$\mathrm{SV}_t \leq \epsilon/3 \leq \epsilon/2$ or $\mathrm{SV}_t > 0$. Note that if we certify $\mathrm{SV}_t \leq \epsilon/3$, then we also know that x_t satisfies

$$\max_{i\in[m]} f^i(x_t, u_t^i) = \max_{y\in Y} \sum_{i=1}^m y^{(i)} f^i(x_t, u_t^i) \leq \epsilon/6 + \inf_{x\in X} \sum_{i=1}^m y_t^{(i)} f^i(x, u_t^i) \leq \epsilon/6 + \mathrm{SV}_t \leq \epsilon/2.$$

Therefore, our point x_t is feasible for (20). Also, when x_t, y_t satisfies (21), we have the bound

$$\begin{aligned}\epsilon^{\bullet}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) &= \max_{y\in Y} \sum_{t=1}^T \theta_t \sum_{i=1}^m y^{(i)} f^i(x_t, u_t^i) - \inf_{x\in X} \sum_{t=1}^T \theta_t \sum_{i=1}^m y_t^{(i)} f^i(x, u_t^i) \\ &\leq \sum_{t=1}^T \theta_t \left[\max_{y\in Y} \sum_{i=1}^m y^{(i)} f^i(x_t, u_t^i) - \inf_{x\in X} \sum_{i=1}^m y_t^{(i)} f^i(x, u_t^i) \right] \leq \epsilon/6.\end{aligned}$$

Consequently, we deduce that the nominal feasibility oracle also naturally bounds $\epsilon^{\bullet}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T)$ although this bound is not utilized in Theorem 4.2. Note that the term $\epsilon^{\bullet}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T)$ inherently includes the functions $\sum_{i=1}^m y^{(i)} f^i(x, u_t^i)$ underlying each SP problem SV_t. At each iteration t, instead of evaluating SV_t to ϵ accuracy, our OFO-based approach performs only a simple update based on the first-order information, and it yields a bound on $\epsilon^{\bullet}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T)$ from the overall collection of these simple updates. This is the motivation for working with the online SP problem. ∎

4.4 Convergence Rates and Discussion

We start by summarizing the convergence rates achievable from our OFO-based approach in various cases. We use the notation $r_u(\epsilon)$ to denote the number of iterations required for algorithms $\mathcal{A}_i$ to generate a sequence $\{u_t\}_{t=1}^T$ such that $\epsilon^{\circ}(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \epsilon/2$. Similarly, we let $r_x(\epsilon)$ be the number of iterations required for $\mathcal{A}_{x,y}$ to generate a sequence $\{x_t, y_t\}_{t=1}^T$ so that $\epsilon^{\bullet}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) \leq \epsilon/2$. Then the resulting worst-case number of iterations needed in Algorithm 1 to obtain robust ϵ-feasibility/infeasibility certificates is $\max\{r_u(\epsilon), r_x(\epsilon)\}$. In order to discuss the total *arithmetic complexity* of each approach, we also let k be the maximum dimension of the uncertain parameters u^i for $i \in [m]$ and n denotes the dimension of the decision variables x.

Under a variety of assumptions on the structure of functions f^i, a careful selection and customization of algorithms $\mathcal{A}_i$, $\mathcal{A}_{x,y}$ and weights θ_t can lead to improved overall convergence rates. For convex uncertainty sets U^i, Table 1 summarizes various assumptions on the structural properties of functions f^i, suggested OFO algorithms from Section 2.3 for the roles of $\mathcal{A}_i$ and $\mathcal{A}_{x,y}$, weights θ_t to be used, along with resulting $r_u(\epsilon), r_x(\epsilon)$ values. Note that the $\log(m)$ terms in the last column result from the set width of Y (see Assumption 2.2 and [20, Remark 3] for further details).

We note that our OFO-based approach returns only robust ϵ-feasible solutions, thus to obtain ϵ-optimal solutions, we need to perform a binary search and repeatedly invoke our method $O(\log(1/\epsilon))$

times. Furthermore, in the case where our domains $X, \{U^i\}_{i=1}^m$ have favorable geometry, such as Euclidean ball or simplex, the vectors $x_t, y_t, \{u_t^i\}_{i=1}^m$ are updated via simple closed-form prox operations, which cost $O(km+mn)$ per iteration. Because the cost of computing the subgradients $\nabla_x f^i(x,u^i), \nabla_u f^i(x,u^i)$ is incurred in each iteration of all of the algorithms, we disregard these in our comparison.

Remark 4.8. The flexibility of our approach extends beyond the cases mentioned in Table 1. Depending on the structure of functions f^i and uncertainty domains U^i, the algorithms $\mathcal{A}_i$ and $\mathcal{A}_{x,y}$ may be replaced by more appropriate OCO algorithms. For example, unless explicitly required by the algorithms $\mathcal{A}_i$, we do not need to assume convexity of the sets U^i. As a result, the follow-the-leader or follow-the-perturbed-leader type algorithms from [27] can be utilized as $\mathcal{A}_i$ in our framework when U^i are nonconvex under certain assumptions ensuring applicability of these algorithms. Such assumptions are satisfied for example when $f^i(x,u^i)$ are linear in u^i and the nonconvex sets U^i admit a certain linear optimization oracle. Similarly, when the functions $f^i(x,u^i)$ are exp-concave in u^i, applying the online Newton step algorithm of [19] for $\mathcal{A}_i$ results in a weighted regret bound of at most $O(\log(T)/T)$ in T iterations. Such f^i that are exp-concave in u^i satisfying Assumption 2.1 arise in optimization under uncertainty problems where variance is used as a risk measure, e.g., mean-variance portfolio optimization problems, see for example [3, Example 25]. ∎

Our modification of the pessimization oracle-based approach of [29] requires $r_x(\epsilon)$ number of calls to the pessimization oracles (19). Specifically, Theorem 4.1 implies that the number of iterations required to solve the robust feasibility problem (4) with pessimization oracles (19) is the number of iterations T needed to guarantee $\epsilon^\bullet(\{x_t,u_t,y_t,\theta_t\}_{t=1}^T) \le \epsilon/2$, given that $\{x_t,y_t\}_{t=1}^T$ are generated in a non-anticipatory way. As indicated in Table 1, this can be done via cheap OFO-based algorithms (see Theorem 2.2) requiring at most $O(\log(m)/\epsilon^2)$ iterations. Taking into account that our binary search factor $O(\log(1/\epsilon))$ to find a robust ϵ-optimal solution, the total number of iterations required is $O(\log(m)\log(1/\epsilon)/\epsilon^2)$, which is much better than the exponential $(1+O(1/\epsilon))^n$ bound of [29, Section 5.2] that uses a full nominal solution oracle (18) as opposed to an algorithm that bounds $\epsilon^\bullet(\{x_t,u_t,y_t,\theta_t\}_{t=1}^T)$. Also, the per-iteration arithmetic cost involves calling m pessimization oracles (19) and a solution oracle. If $\sup_{u^i\in U^i} f^i(x,u^i)$ has a simple closed form solution, then the resulting arithmetic cost is $O(k)$ for each pessimization oracle. If we can use polynomial-time IPMs, the cost becomes $O(k^3\log(1/\epsilon))$ (see [7, Section 6.6]), and using FOMs has cost $O(k\log(1/\epsilon))$ in the best case when the f^i are smooth and strongly convex in u^i.

The approach of [5] requires $r_u(\epsilon)$ number of calls to the nominal optimization oracle (or $r_u(\epsilon)\log(1/\epsilon)$ calls to the nominal feasibility oracle). Then their overall convergence rate in terms of basic arithmetic operations depends on the type of solver used to solve the nominal optimization/feasibility problem (20). We next examine the cost of implementing these oracles.

When applicable, polynomial-time IPMs are guaranteed to terminate in $O(\sqrt{m}\log(1/\epsilon))$ iterations with a solution to (20) and thus offer the best rates in terms of their dependence on ϵ. They also have the advantage that they can act as a nominal optimization oracle, and hence by Remark 4.6 there will be no need to perform an additional binary search to find an ϵ-optimal solution. On the other hand, they demand significantly more memory and their per-iteration cost is quite high in terms of the dimension, usually around the order of $O(n^3+mn)$, see [7, Chapter 6.6]. In order to keep both the memory requirements and the per-iteration cost associated with implementing the nominal feasibility oracle low, we may opt for an FOM called the CoMirror algorithm, see [1] and [25, Section 1.3]. The CoMirror algorithm is guaranteed to find a solution to the nominal ϵ-feasibility problem within $O(1/\epsilon^2)$ iterations, with a much cheaper per-iteration

Table 2: Summary of different approaches to generate $\{x_t, u_t, y_t\}_{t=1}^T$.

Approach	Binary search	No. iterations	Per-iteration cost
OFO-based	$\log(1/\epsilon)$	$\max\{r_u(\epsilon), r_x(\epsilon)\}$	$O(km+mn)$
Closed form pessimization oracle	$\log(1/\epsilon)$	$r_x(\epsilon)$	$O(km+mn)$
IPM pessimization oracle	$\log(1/\epsilon)$	$r_x(\epsilon)$	$O(k^3 m\log(1/\epsilon)+mn)$
FOM pessimization oracle (when f^i are smooth/str. conv. in u^i)	$\log(1/\epsilon)$	$r_x(\epsilon)$	$O\left(km\log(1/\epsilon)+mn\right)$
IPM feasibility oracle	1	$r_u(\epsilon)$	$O\left(km+\sqrt{m}(n^3+mn)\log(1/\epsilon)\right)$
CoMirror feasibility oracle	1	$r_u(\epsilon)$	$O\left(km+mn/\epsilon^2\right)$
Convex-concave SP feas. oracle (when f^i are smooth in x)	$\log(1/\epsilon)$	$r_u(\epsilon)$	$O\left(km+\log(m)mn/\epsilon\right)$

cost of $O(mn)$. Because the CoMirror method can optimize as well, there is no need to perform binary search. However, it cannot exploit further structural properties of the functions f^i, such as smoothness in x, to improve the dependence on ϵ. In order to exploit such properties, it is possible to cast (20) as a convex-concave SP problem as in Remark 4.7, and then apply efficient FOMs such as Nesterov's algorithm [32] or Nemirovski's Mirror Prox algorithm [30] to achieve a convergence rate of $O(\log(m)/\epsilon)$. Both of these FOMs have the per-iteration cost of $O(mn)$. This convex-concave SP approach can only be used as a nominal feasibility oracle, so we must repeat the process $\log(1/\epsilon)$ times to obtain an ϵ-optimal solution.

In Table 2, we summarize the rates for the various approaches. Note that the total *arithmetic complexity* of each approach is obtained by multiplying the quantities in each row. Furthermore, the quantities $r_u(\epsilon), r_x(\epsilon)$ are given Table 1.

Table 2 indicates that our modified pessimization oracle-based approach when it admits a closed form solution and the nominal feasibility oracle-based approach which uses a polynomial-time IPM solver to implement the oracle give the best dependence on ϵ among all of the methods. These are better than our OFO-based approach by factors of $\max\{1, r_u(\epsilon)/r_x(\epsilon)\}$ and $\max\{1, r_x(\epsilon)/r_u(\epsilon)\}$ respectively. However, unlike our OFO-based approach, these methods are restricted to using non-anticipatory algorithms to generate the iterates $[x_t; y_t]$ and u_t respectively, which prevents possible acceleration. Furthermore, in many applications, we can expect that $r_u(\epsilon) \approx r_x(\epsilon)$. In this case, our OFO-based approach becomes competitive with having a closed form pessimization oracle or using a nominal IPM solver in [5]. However, compared to IPMs, it is able to maintain a much lower dependence on the dimensions m, n and thus is much more scalable, whereas the cost per iteration of such IPMs has a rather high dependence on the dimension. In addition, their memory requirements are far more than OFO algorithms, posing a critical disadvantage to their use in large-scale applications. Similar comparisons of our OFO-based approach against pessimization or nominal oracle-based approaches utilizing other methods point out its advantage, which is at least an order of magnitude better in terms of its dependence on ϵ. In fact, when $r_x(\epsilon) \approx r_u(\epsilon)$, our method can lead to savings over the approach of [5] with CoMirror algorithm used in its oracle by a factor as large as $O(1/(\epsilon^2\log(1/\epsilon)))$.

Remark 4.9. Ben-Tal et al. [5, Section 5] pose the following open question: are $\Omega(1/\epsilon^2)$ nominal feasibility oracle calls indeed necessary to solve the robust feasibility problem (4) with their approach? When the functions f^i are strongly concave in u^i, our analysis summarized in Table 1 leads to a rate of $r_u(\epsilon) = O(1/\epsilon)$. Therefore, we deduce from Table 2 that under this strong concavity assumption on f^i, we only need $O(1/\epsilon)$ nominal feasibility/optimization oracle calls to solve (4). Thus, this lower bound of $\Omega(1/\epsilon^2)$ does not hold under favorable circumstances. Nonetheless, our results only partially address this open question because they do not lead to any conclusions in the general case, i.e., without favorable structural assumptions on f^i. Indeed, even under the strong concavity assumption, establishing a matching lower bound on $r_u(\epsilon)$ remains open. ∎

Finally, let us get back to the case when we have a robust feasibility problem with a single constraint $m = 1$ and convex uncertainty set $U = U^1$. In such a case, as discussed in Remark 3.1, we have a direct convex-concave SP problem under Assumption 2.1; without any further structural assumptions on f^1 such a problem can be solved in $O(1/\epsilon^2)$ iterations. Note that our approach also achieves this rate immediately, cf. Tables 1 and 2. Moreover, when f^1 is smooth in x and strongly concave in u^1, our general framework can achieve a rate of $O(1/\epsilon)$. However, directly working with the convex-concave SP formulation can improve this rate further. For example, such an improved rate of $O(1/\sqrt{\epsilon})$ is achieved in [2] for a robust support vector machine problem.

5 Conclusion

In this paper, we advance the line of research in [12, 13, 29, 5] that aims to solve robust optimization problems via iterative techniques, i.e., without transforming them into their equivalent robust counterparts. Thus far, the prior literature on iterative methods for RO rely on more expensive pessimization or nominal feasibility oracles. However, in many applications of robust convex optimization, the original deterministic problem comes equipped with first-order oracles capable of providing gradient/subgradient information on the constraint functions f^i. In this paper, we present an efficient framework that can work with both cheap online first-order oracles, but also captures the prior oracle-based approaches of [29] and [5]. We further show that working with these OFO oracles essentially does not increase the number of overall oracle calls, i.e., the number of main iterations of our approach is better than or comparable to the prior approaches. Moreover, when OFO oracles are utilized in our framework, the resulting overall arithmetic complexity including all of the basic operations in each iteration is remarkably cheaper than the prior approaches. The resulting framework is simple, easy-to-implement, flexible, and it can easily be customized to many applications. We discuss how our framework works and compares to [5] and also demonstrate this via illustrative examples in robust SVM and robust quadratic programming. For example, in the context of robust QPs, the most expensive operation in each iteration of our framework is maximum eigenvalue computations.

Our work partially resolves an open question stated in [5] for the lower bound on the number of iterations/calls needed in their nominal feasibility oracle based framework. Specifically, [5] poses the question of whether their upper bound of $O(1/\epsilon^2)$ iterations/oracle calls to obtain a robust ϵ-feasibility/infeasibility certificate is the best possible. We show that when the constraint functions possess favorable structure such as strong concavity, by exploiting this structural information and using the customization of weighted regret online mirror descent algorithm for strongly convex functions given in [20, Theorem 2], it is possible to achieve a better convergence rate of $O(1/\epsilon)$ in both the nominal feasibility oracle framework of [5] and our online first-order oracle setup. However, it remains open whether $O(1/\epsilon^2)$ bound is tight when no further favorable structure is present in

f^i. Likewise, the tightness of $O(1/\epsilon)$ bound in the case of strongly concave functions f^i is also open.

Our results also lead to several other compelling avenues for future research. From a practical perspective, it is well-known, and also confirmed by our preliminary proof-of-concept computational experiments, that the computation of gradients/subgradients constitute a major bottleneck in the practical performance of FOMs. Thus, as a step to reduce the efforts involved in subgradient computations, possible incorporation of stochastic [34, 31] and/or randomized FOMs [24, 10] working with stochastic subgradients into our framework is of great practical and theoretical interest. A critical assumption in our approach as well as others, e.g., see [5] and references therein, is that the domain X is convex. Removing the convexity requirement on the domain X will be an important theoretical development on its own. Besides, this will open up possibilities for more principled approaches to solving robust combinatorial optimization problems (see [11, 12]) where such a convexity assumption on X is not satisfied. Finally, another compelling research direction is develop analogous frameworks for multi-stage RO problems such as robust Markov decision processes (see [33, 22]).

Acknowledgments

This research is supported in part by NSF grant CMMI 1454548.

The authors would like to thank Arkadi Nemirovski for suggesting the convex-concave reformulation of the robust feasibility problem presented in Section A, in particular Lemma A.2.

References

[1] A. Beck, A. Ben-Tal, N. Guttmann-Beck, and L. Tetruashvili. The CoMirror algorithm for solving nonsmooth constrained convex problems. *Operations Research Letters*, 38(6):493 – 498, 2010.

[2] A. Ben-Tal, S. Bhadra, C. Bhattacharyya, and A. Nemirovski. Efficient methods for robust classification under uncertainty in kernel matrices. *Journal of Machine Learning Research*, 13(1):2923–2954, Oct. 2012.

[3] A. Ben-Tal, D. den Hertog, and J.-P. Vial. Deriving robust counterparts of nonlinear uncertain inequalities. *Mathematical Programming*, 149(1-2):265–299, 2015.

[4] A. Ben-Tal, L. Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton University Press, 2009.

[5] A. Ben-Tal, E. Hazan, T. Koren, and S. Mannor. Oracle-based robust optimization via online learning. *Operations Research*, 63(3):628–638, 2015.

[6] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998.

[7] A. Ben-Tal and A. Nemirovski. *Lectures on Modern Convex Optimization*. Society for Industrial and Applied Mathematics, 2001.

[8] A. Ben-Tal and A. Nemirovski. Robust optimization – methodology and applications. *Mathematical Programming*, 92(3):453–480, 2002.

[9] A. Ben-Tal and A. Nemirovski. Selected topics in robust convex optimization. *Mathematical Programming*, 112(1):125–158, 2008.

[10] A. Ben-Tal and A. Nemirovski. On solving large-scale polynomial convex problems by randomized first-order algorithms. *Mathematics of Operations Research*, 40(2):474–494, 2015.

[11] D. Bertsimas, D. B. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.

[12] D. Bertsimas and M. Sim. Robust discrete optimization and network flows. *Mathematical Programming*, 98(1-3):49–71, 2003.

[13] G. Calafiore and M. Campi. Uncertain convex programs: randomized solutions and confidence levels. *Mathematical Programming*, 102(1):25–46, 2005.

[14] C. Caramanis, S. Mannor, and H. Xu. Robust optimization in machine learning. In S. Sra, S. Nowozin, and S. Wright, editors, *Optimization for Machine Learning*. MIT Press, 2012.

[15] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

[16] L. El Ghaoui and H. Lebret. Robust solutions to least-squares problems with uncertain data. *SIAM Journal on Matrix Analysis and Applications*, 18(4):1035–1064, 1997.

[17] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins Studies in Mathematical Sciences. The Johns Hopkins University Press, Baltimore, MD, 1996.

[18] E. Hazan. The convex optimization approach to regret minimization. In S. Sra, S. Nowozin, and S. Wright, editors, *Optimization for Machine Learning*. MIT Press, 2012.

[19] E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.

[20] N. Ho-Nguyen and F. Kılınç-Karzan. The role of flexibility in structure-based acceleration for online convex optimization. Technical report, August 2016. http://www.optimization-online.org/DB_HTML/2016/08/5571.html.

[21] N. Ho-Nguyen and F. Kılınç-Karzan. A second-order cone based approach for solving the trust region subproblem and its variants. Technical report, March 2016. http://www.optimization-online.org/DB_HTML/2016/03/5362.html.

[22] G. N. Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280, 2005.

[23] T. Joachims. Training linear svms in linear time. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 217–226. ACM, 2006.

[24] A. Juditsky, F. Kılınç-Karzan, and A. Nemirovski. Randomized first order algorithms with applications to ℓ_1 minimization. *Mathematical Programming*, 142(1-2):269–310, 2013.

[25] A. Juditsky and A. Nemirovski. First-order methods for nonsmooth convex large-scale optimization, I: General purpose methods. In S. Sra, S. Nowozin, and S. Wright, editors, *Optimization for Machine Learning*. MIT Press, 2012.

[26] A. Juditsky and A. Nemirovski. First-order methods for nonsmooth convex large-scale optimization, II: Utilizing problem's structure. In S. Sra, S. Nowozin, and S. Wright, editors, *Optimization for Machine Learning*. MIT Press, 2012.

[27] A. Kalai and S. Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291 – 307, October 2005.

[28] Q. V. Le, A. J. Smola, and S. Vishwanathan. Bundle methods for machine learning. In *Advances in Neural Information Processing Systems*, pages 1377–1384, 2007.

[29] A. Mutapcic and S. Boyd. Cutting-set methods for robust convex optimization with pessimizing oracles. *Optimization Methods and Software*, 24(3):381–406, June 2009.

[30] A. Nemirovski. Prox-method with rate of convergence $O(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.

[31] A. S. Nemirovski, A. B. Juditsky, G. Lan, and A. Shapiro. Stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.

[32] Y. Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.

[33] A. Nilim and L. El Ghaoui. Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.

[34] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407, 1951.

[35] S. Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.

[36] P. K. Shivaswamy, C. Bhattacharyya, and A. J. Smola. Second order cone programming approaches for handling missing and uncertain data. *Journal of Machine Learning Research*, 7(Jul):1283–1314, 2006.

[37] V. Vapnik. *The Nature of Statistical Learning Theory*. Information Science and Statistics. Springer-Verlag New York, 2000.

[38] H. Xu, C. Caramanis, and S. Mannor. Robustness and regularization of support vector machines. *Journal of Machine Learning Research*, 10(Jul):1485–1510, 2009.

[39] H. Xu, C. Caramanis, and S. Mannor. Robust regression and Lasso. *IEEE Transactions on Information Theory*, 56(7):3561–3574, July 2010.

A Convex-Concave Saddle Point Reformulation

The SP problem (8) based on the function $\Phi(x,u)$ which is not necessarily concave in u admits a convex-concave SP representation in a lifted space via perspective transformations. To present this reformulation, we start by defining the following sets with additional variables $y \in \mathbb{R}^m_+$ and new variables v^i for $i \in [m]$:

$$V^i = \left\{ [v^i; y^{(i)}] : \ 0 < y^{(i)} \leq 1, \ \frac{v^i}{y^{(i)}} \in U^i \right\} \quad \forall i \in [m],$$
$$W = \left\{ w = [v^1; \ldots; v^m; y] : \ [v^i; y^{(i)}] \in \mathrm{cl}(V^i), \ i \in [m], \ \sum_{i=1}^m y^{(i)} = 1 \right\}.$$

Note that for all $i \in [m]$, $\mathrm{cl}(V^i) = V^i \cup \{[0;0]\}$ because we assumed U^i to be closed sets. For the point $[v^i; y^{(i)}] = [0;0]$, we set $y^{(i)} f^i \left(x, \frac{v^i}{y^{(i)}} \right) = 0$ for any $x \in X$. Note that setting $y^{(i)} f^i \left(x, \frac{v^i}{y^{(i)}} \right) = 0$ for $[v^i; y^{(i)}] = [0;0]$ is well-defined as the continuation since from Assumption 2.1, $f^i(x,u^i)$ is continuous and finite-valued on U^i, and U^i is compact, so we can deduce that $f^i(x,u^i)$ will be bounded on U^i.

We also define the function $\psi : X \times W \to \mathbb{R}$ as

$$\psi(x,w) = \psi(x,v,y) := \sum_{i=1}^m y^{(i)} f^i \left(x, \frac{v^i}{y^{(i)}} \right).$$

Lemma A.1. *For fixed $w \in W$, the function $\psi(x,w)$ is convex in x over X, and $\psi(x,w)$ is a concave function of w over W for any fixed x. Moreover, W is closed, and when U^i for $i \in [m]$ are convex, the sets V^i for $i \in [m]$ and W are all convex.*

Proof. For any $w = [v^1; \ldots; v^m; y]$, the function ψ is convex in x since in all of the nonzero terms in the summation over all $i \in [m]$ defining ψ, we have $y^{(i)} > 0$ and in each such nonzero term each function $f^i \left(x, \frac{v^i}{y^{(i)}} \right)$ is convex in x for the given $\frac{v^i}{y^{(i)}} \in U^i$ (see Assumption 2.1). In addition, for any given $x \in X$, the function ψ is jointly concave in v and y because it is written as a sum of the perspective functions of functions f^i which are concave in u^i (see Assumption 2.1).

The closedness of W is immediate, and the convexity of the sets V^i and W follows immediately from their definition and the convexity assumption on U^i. □

With these definitions and Lemma A.1, we observe that (8) is equivalent to evaluating the convex-concave SP problem defined by the function ψ over the convex domains X and W:

$$\inf_{x \in X} \sup_{w \in W} \psi(x,w) \leq \epsilon \quad \text{or} \quad \inf_{x \in X} \sup_{w \in W} \psi(x,w) > 0. \tag{22}$$

We state this formally in the following lemma.

Lemma A.2. *For any $\epsilon > 0$ and $\bar{x} \in X$,*

$$\max_{i \in [m]} \sup_{u^i \in U^i} f^i(\bar{x}, u^i) \leq \epsilon \quad \textit{if and only if} \quad \sup_{w \in W} \psi(\bar{x}, w) \leq \epsilon.$$

As a result,

$$\inf_{x \in X} \max_{i \in [m]} \sup_{u^i \in U^i} f^i(x, u^i) \leq \epsilon \quad \textit{if and only if} \quad \inf_{x \in X} \sup_{w \in W} \psi(x, w) \leq \epsilon.$$

Proof. Fix $\bar{x} \in X$ and $\epsilon > 0$. Suppose $\max_{i\in[m]} \sup_{u^i\in U^i} f^i(\bar{x}, u^i) \le \epsilon$; then for all $u^i \in U^i$, $i \in [m]$, we have $f^i(\bar{x}, u^i) \le \epsilon$. Now consider any $w = [v^1; \ldots; v^m; y] \in W$. Then $0 \le y^{(i)} \le 1$ for all $i \in [m]$ and $\sum_{i=1}^m y^{(i)} = 1$. For all $i \in [m]$, define $u^i = \frac{v^i}{y^{(i)}} \in U^i$ whenever $y^{(i)} > 0$. Then $y^{(i)} f^i(\bar{x}, \frac{v^i}{y^{(i)}}) = y^{(i)} f^i(\bar{x}, u^i) \le y^{(i)}\epsilon$ for $0 < y^{(i)} \le 1$. In addition, when $y^{(i)} = 0$, because $w \in W$ we must have $v^i = 0$ and then by definition we have $y^{(i)} f^i(\bar{x}, \frac{v^i}{y^{(i)}}) = 0$. Therefore, from $\sum_{i=1}^m y^{(i)} = 1$, we deduce $\psi(\bar{x}, w) = \sum_{i=1}^m y^{(i)} f^i\left(\bar{x}, \frac{v^i}{y^{(i)}}\right) \le \epsilon$ holds for any $w \in W$.

Now suppose that $\sup_{w\in W} \psi(\bar{x}, w) \le \epsilon$ holds. Given $i \in [m]$ and $u^i \in U^i$, set w to have components $y^{(i)} = 1$, $v^i = u^i$, and $[v^j; y^{(j)}] = [0; 0]$ for $j \ne i$. Then $f^i(\bar{x}, u^i) = \psi(\bar{x}, w) \le \epsilon$. Hence, $\max_{i\in[m]} \sup_{u^i\in U^i} f^i(\bar{x}, u^i) \le \epsilon$ follows. □

Remark A.1. When $m = 1$, i.e., we have only one function $f^1(x, u^1)$ and only one uncertainty set U^1, hence $W = U^1$ and $\inf_{x\in X} \sup_{v\in W} \psi(x, w) = \inf_{x\in X} \sup_{u^1\in U^1} f^1(x, u^1)$. Also, under Assumption 2.1, $\psi(x, w)$ is convex in x and concave in u^1. Thus, the preceding perspective transformation resulting in (22) directly generalizes this case of a convex-concave SP formulation for $m = 1$ discussed in Remark 3.1. ∎

As a result, Lemma A.2 and Theorem 3.1 combined with any FOM that provides bounds on the saddle point gap $\epsilon^{\psi}_{\text{sad}}(\bar{x}, \bar{w})$ lead to an efficient way of verifying robust feasibility of (8) as follows:

Theorem A.1. *Suppose $\bar{x} \in X$, $\bar{w} \in W$, and $\tau \in (0, 1)$ are such that $\epsilon^{\psi}_{\text{sad}}(\bar{x}, \bar{w}) \le \tau\epsilon$. If $\psi(\bar{x}, \bar{w}) \le (1 - \tau)\epsilon$, then $\max_{i\in[m]} \sup_{u^i\in U^i} f^i(\bar{x}, u^i) \le \epsilon$. If $\psi(\bar{x}, \bar{w}) > (1 - \tau)\epsilon$ and $\tau \le \frac{1}{2}$, then $\inf_{x\in X} \max_{i\in[m]} \sup_{u^i\in U^i} f^i(x, u^i) > 0$.*

Proof. Suppose $\psi(\bar{x}, \bar{w}) \le \tau\epsilon$. By Theorem 3.1, we have $\inf_{x\in X} \sup_{w\in W} \psi(x, w) \le \sup_{w\in W} \psi(\bar{x}, w) \le \epsilon$. By Lemma A.2, $\max_{i\in[m]} \sup_{u^i\in U^i} f^i(\bar{x}, u^i) \le \epsilon$ as well.

On the other hand, when $\psi(\bar{x}, \bar{w}) > (1 - \tau)\epsilon$ and $\tau \le \frac{1}{2}$, Theorem 3.1 implies $\inf_{x\in X} \sup_{w\in W} \psi(x, w) \ge \inf_{x\in X} \psi(x, \bar{w}) > 0$. Then by Lemma A.2, $\max_{i\in[m]} \sup_{u^i\in U^i} f^i(\bar{x}, u^i) > 0$ follows. □

Because of the existence of efficient FOMs to solve convex-concave SP problems, Theorem A.1 suggests a possible advantage of using the convex-concave SP problem given in (22). Nevertheless, working with the SP reformulation given by (22) in the extended space $X \times W$ presents a number of critical challenges. First, efficient FOMs associated with convex-concave SP problems often require computing prox operations or projections onto the domains X and W. Unfortunately, even if projection (or prox-mappings) onto U^i admits a closed form solution or an efficient procedure, it is unclear how to extend such projections onto W. Furthermore, while the perspective transformations involved in constructing the function ψ preserves certain desirable properties of the functions f^i, such as Lipschitz continuity and smoothness, the parameters associated with ψ are in general larger than those associated with the original functions f^i. Such parameters are critical for FOM convergence rates, and thus the FOMs when applied to solve (22) will have slower convergence rates.

To address the issues outlined above, in the main paper we discuss how to obtain robust feasibility/infeasibility certificates for the convex-nonconcave SP problem (8) directly, i.e., we work with the functions f^i and the sets U^i directly. This direct approach in particular allows us to take greater advantage of the structure of the original formulation such as the availability of efficient

projection (prox) computations over domains, and/or better parameters for smoothness, Lipschitz continuity, etc., of the functions.

B Applications and Examples

In this section we discuss the setup and resulting convergence rates for our framework for robust support vector machine problems and for robust quadratic programs.

B.1 Robust Support Vector Machines

Consider the linear classification problem: given labelled training data $\{a^i, b^i\}_{i=1}^m$ where $a^i \in \mathbb{R}^n$ represents the features of the data and $b^i \in \{\pm 1\}$ corresponds to its label, we wish to construct a linear classifier $h : \mathbb{R}^n \to \{\pm 1\}$ of the form $h(a) = \text{sign}(\langle x, a\rangle)$. Support vector machines (SVMs) generate $x \in \mathbb{R}^n$ by minimizing the hinge loss $l(x, a^i, b^i) = \max\{0, 1 - b^i\langle x, a^i\rangle\}$ on the training data with ℓ_2-norm regularization, which can be cast as the following problem (see [37, Chapter 5.5.1])

$$\min_{x,\zeta} \left\{ \sum_{i=1}^m \zeta^i : \begin{array}{ll} \zeta^i \geq 1 - b^i\langle x, a^i\rangle, & i \in [m] \\ \zeta^i \geq 0, & i \in [m] \\ \|x\|_2 \leq R & \end{array} \right\}.$$

Here, the constant $R > 0$ represents the level of regularization desired. Suppose now that our data vectors a^i are subject to uncertainty. Then a robust formulation of SVM with uncoupled ellipsoidal uncertainty [36, Section 2.1] is

$$\min_{x,\zeta} \left\{ \sum_{i=1}^m \zeta^i : \begin{array}{ll} 0 \geq \sup_{u^i \in U^i} \left\{1 - \zeta^i - b^i\langle x, a^i + u^i\rangle\right\}, & i \in [m] \\ \zeta^i \geq 0, & i \in [m] \\ \|x\|_2 \leq R & \end{array} \right\}, \tag{23}$$

where the uncertainty sets are given by Euclidean norm balls $U^i = \left\{u^i : \|u^i\|_2 \leq 1\right\}$. Then our constraint functions are

$$f^i(x, \zeta, u^i) = 1 - \zeta^i - b^i\langle x, a^i + u^i\rangle, \quad \text{for } i \in [m] \tag{24}$$

which are convex in $[x; \zeta]$ and concave in u^i. From conic duality, we see that robust counterpart of this problem is a convex QP. While convex QPs can be solved efficiently via polynomial-time IPMs, scalability of the resulting problem becomes an issue especially in the case of classification problems with millions of data points.

As in Section 4.3, Ben-Tal et al. [5] address this scalability issue by applying regret-minimizing OCO algorithms such as subgradient descent to generate a realization of the noise $\{u_{t+1}^i\}_{i=1}^m$, then employ a nominal feasibility oracle to obtain x_{t+1} and ζ_{t+1}. More precisely, they set:

$$u_{t+1}^i = \text{Proj}(u_t^i + \gamma_t \nabla_{u^i} f^i(x_t, \zeta_t, u_t^i)) = \text{Proj}(u_t^i + \gamma_t b^i x_t) \quad \forall i \in [m],$$

where Proj denotes projection onto the unit ball, and their oracle solves the feasibility problem

$$\begin{cases} \textit{Either}\text{: find } [x; \zeta] \in X \times Z \quad \text{s.t.} \quad 1 - \zeta^i - b^i\langle x, a^i + u_t^i\rangle \leq \epsilon/2 \quad \forall i \in [m]; \\ \textit{or}\text{: declare infeasibility, } \forall [x; \zeta] \in X \times Z, \ \exists i \in [m], \quad \text{s.t.} \quad 1 - \zeta^i - b^i\langle x, a^i + u^i\rangle > 0, \end{cases} \tag{25}$$

where $X = \{x : \|x\|_2 \leq R\}$ and $Z = \left\{\zeta : \sum_{i=1}^m \zeta^i \leq C, \ \zeta \geq 0\right\}$. Here, the constraint $\sum_{i=1}^m \zeta^i \leq C$ is the result of embedding the objective into the constraints, which allows us to minimize the robust

SVM problem by solving a sequence of convex QP feasibility problems via the nominal feasibility oracle. Alternatively, since nominal SVM problems are quite well-studied, we could use a nominal optimization oracle to solve

$$\mathrm{Opt}_t = \inf_{[x;\zeta]} \left\{ f^0(x) : 1 - \zeta^i - b^i \langle x, a^i + u_t^i \rangle \le 0,\ i \in [m], \quad [x;\zeta] \in X \times \zeta \right\} \tag{26}$$

and, by Remark 4.6, avoid having to perform a binary search to optimize. In general, to achieve ϵ-optimal ϵ-feasible solution to (26), most polynomial-time IPMs require $O(m^2 n \log(1/\epsilon))$ runtime. However, the dependence m^2 on the number of training samples makes these inappropriate for large-scale applications. An improvement to this is the SVM-Perf algorithm of Joachims [23], where [28] demonstrate convergence within $O(mn/\epsilon)$ runtime. Since the constraint functions f^i in (24) are linear in u^i, OCO algorithms to bound $\epsilon^\circ(\{[x_t;\zeta_t], u_t, \theta_t\}_{t=1}^T) \le \epsilon/2$ require at least $r_u(\epsilon) = O(1/\epsilon^2)$ iterations T (see Table 1). Each iteration solves a single instance of (26), and also has running time $O(mn)$ to update u_t, therefore the total running time for the oracle-based approach of [5] to solve the robust SVM problem (23) is $O((mn + mn/\epsilon)/\epsilon^2) = O(mn/\epsilon^3)$.

Our OFO-based approach detailed in Section 4.1 also generates u_t in the same way, and thus also requires $r_u(\epsilon) = O(1/\epsilon^2)$ iterations to bound $\epsilon^\circ(\{[x_t;\zeta_t], u_t, \theta_t\}_{t=1}^T) \le \epsilon/2$. But in contrast, we generate x_t and ζ_t by solving an online SP problem with the sequence of functions

$$\phi_t(x, \zeta, y) = \sum_{i=1}^m y^i (1 - \zeta^i - b^i \langle x, a^i + u_t^i \rangle)$$

over the domains $[x;\zeta] = X \times Z$ and $y \in Y = \Delta_m$. Then [20, Theorem 3] using the online subgradient/mirror descent algorithm leads to the updates of our variables as

$$\begin{aligned}
x_{t+1} &= \mathrm{Proj}\left(x_t - \gamma_t \nabla_x \phi_t(x_t, \zeta_t, y_t)\right) = \mathrm{Proj}\left(x_t + \gamma_t \sum_{i=1}^m y_t^i b^i (a^i + u_t^i)\right) \\
\zeta_{t+1} &= C\,\mathrm{Prox}_{\zeta_t/C}\left(C\gamma_t \nabla_\zeta \phi_t(x_t, \zeta_t, y_t)\right) = C\,\mathrm{Prox}_{\zeta_t/C}\left(C\gamma_t y_t\right) \\
y_{t+1} &= \mathrm{Prox}_{y_t}\left(-\gamma_t \nabla_y \phi_t(x_t, \zeta_t, y_t)\right) = \mathrm{Prox}_{y_t}\left(-\gamma_t \left\{\zeta_t^i + b^i \langle x_t, a^i + u_t^i \rangle\right\}_{i=1}^m\right),
\end{aligned}$$

where Prox is the usual prox-mapping onto the simplex. The number of iterations T to guarantee $\epsilon^\bullet(\{[x_t;\zeta_t], u_t, y_t, \theta_t\}_{t=1}^T) \le \epsilon/2$ is $r_x(\epsilon) = O(\log(m)/\epsilon^2)$. Also, each iteration takes time $O(mn)$ to update the variables u_t, x_t, ζ_t, y_t. Furthermore, with our OFO-based approach, we must perform a binary search to find a robust ϵ-optimal solution to (23), so we must repeat solving the robust feasibility problem $O(\log(1/\epsilon))$ times. Then the total running time for our method is $O(mn \max\{r_u(\epsilon), r_x(\epsilon)\} \log(1/\epsilon)) = O(m \log(m) n \log(1/\epsilon)/\epsilon^2)$. By employing our OFO-based approach as opposed to the oracle-based approach of [5], we improve the accuracy dependence $\log(1/\epsilon)/\epsilon^2$ as opposed to $1/\epsilon^3$, but worsen the dimension dependence by a factor $\log(m)$. In most cases, this is a very good trade-off because $\log(m)$ will be quite small even for large m.

B.2 Robust Quadratic Programming

We next consider a robust feasibility problem of a quadratically constrained quadratic program (QP) with ellipsoidal uncertainty. To be precise, our deterministic feasibility problem is

$$\text{find} \quad x \in X \quad \text{s.t.} \quad \|A_i x\|_2^2 \le b_i^\top x + c_i, \quad \forall i \in [m],$$

where $X \subseteq \mathbb{R}^n$ is the unit Euclidean ball, $A_i \in \mathbb{R}^{n\times n}$, $b_i \in \mathbb{R}^n$, and $c_i \in \mathbb{R}$ for all $i \in [m]$. We consider the robust quadratic feasibility problem given by

$$\text{find} \quad x \in X \quad \text{s.t.} \quad \sup_{u\in\widehat{U}} \left\| \left(A_i + \sum_{k=1}^{K} P_k^i \, u^{(k)} \right) x \right\|_2^2 - b_i^\top x - c_i \leq 0, \quad \forall i \in [m], \tag{27}$$

where $P_1^i, \ldots, P_K^i$ are uncertainty matrices for each constraint $i \in [m]$, for simplicity we assume uncertainty sets $U^i = \widehat{U} = \{u \in \mathbb{R}^K : \|u\|_2 \leq 1\}$ for all $i \in [m]$, and $u^{(k)}$ denotes the k-th entry of u.

It is well known that the robust counterpart of this feasibility problem is a semidefinite program [4, 11]. Because current state-of-the-art QP solvers can handle two to three orders of magnitude larger QPs than semidefinite programs (SDPs), Ben-Tal et al. [5, Section 4.2] suggest an approach that avoids solving SDPs associated with robust QPs. Their approach relies on running a probabilistic OCO algorithm in which a trust region subproblem (TRS)—a class of well-studied nonconvex QPs—is solved in each iteration. Our results here further enhance this approach. In particular, we show that we can achieve the same rate of convergence in our framework while working with a deterministic OCO algorithm and only carrying out first-order updates in each iteration. In fact, the most expensive operation involved with each iteration of our approach is a maximum eigenvalue computation. Because maximum eigenvalue computation is much cheaper than solving a TRS, we not only present a deterministic approach but also strikingly reduce the cost of each iteration.

To simplify our exposition, let us introduce some notation. For each $i \in [m]$, we define the matrix $\mathcal{P}_x^i \in \mathbb{R}^{n\times K}$ whose columns are given by the vectors $P_k^i x$ for $k \in [K]$ together with

$$Q_x^i := (\mathcal{P}_x^i)^\top \mathcal{P}_x^i \in \mathbb{S}_+^K, \quad r_x^i := (\mathcal{P}_x^i)^\top A_i x \in \mathbb{R}^K, \quad \text{and} \quad s_x^i := \|A_i x\|_2^2 - b_i^\top x - c_i \in \mathbb{R};$$

then it is easy to check that for all $i \in [m]$ and $u \in \mathbb{R}^K$ we have

$$\left\| \left(A_i + \sum_{k=1}^{K} P_k^i \, u^{(k)} \right) x \right\|_2^2 - b_i^\top x - c_i = u^\top Q_x^i u + 2(r_x^i)^\top u + s_x^i.$$

For each $i \in [m]$, we define $f^i : X \times \widehat{U} \to \mathbb{R}$ as

$$\begin{aligned} f^i(x,u) &:= \left\| \left(A_i + \sum_{k=1}^{K} P_k^i \, u^{(k)} \right) x \right\|_2^2 - b_i^\top x - c_i + \lambda_{\max}(Q_x^i) \left(1 - \|u\|_2^2\right) \\ &= u^\top Q_x^i u + 2(r_x^i)^\top u + s_x^i + \lambda_{\max}(Q_x^i) \left(1 - \|u\|_2^2\right). \end{aligned} \tag{28}$$

Lemma B.1. *For each $i \in [m]$, the function $f^i(x,u)$ defined in (28) is convex in x for any fixed $u \in \widehat{U}$ and concave in u for any given x. Moreover, for all $i \in [m]$ and for any $x \in X$,*

$$\sup_{u\in\widehat{U}} \left\| \left(A_i + \sum_{k=1}^{K} P_k^i \, u^{(k)} \right) x \right\|_2^2 - b_i^\top x - c_i = \sup_{u\in\widehat{U}} f^i(x,u).$$

Proof. Fix $i \in [m]$. By rearranging terms in (28), we obtain $f^i(x,u) = u^\top (Q_x^i - \lambda_{\max}(Q_x^i) I_K) u + 2(r_x^i)^\top u + s_x^i$. Since $Q_x^i - \lambda_{\max}(Q_x^i) I_K \in \mathbb{S}_+^K$ for any given x, $f^i(x,u)$ is concave in u for any given x.

Now consider a fixed $u \in \widehat{U}$. Note that

$$\lambda_{\max}(Q_x^i) = \max_{\|v\|_2 \le 1} v^\top (Q_x^i) v = \max_{\|v\|_2 \le 1} \sum_{1 \le j,k \le K} v^{(j)} v^{(k)} x^\top (P_j^i)^\top P_k^i \, x$$
$$= \max_{\|v\|_2 \le 1} x^\top \left(\sum_{k=1}^K P_k^i \, v^{(k)} \right)^\top \left(\sum_{k=1}^K P_k^i \, v^{(k)} \right) x.$$

Because $\left(\sum_{k=1}^K P_k^i \, v^{(k)} \right)^\top \left(\sum_{k=1}^K P_k^i \, v^{(k)} \right) \in \mathbb{S}_+^n$, then $\lambda_{\max}(Q_x^i)$ is a maximum of convex quadratic functions of x and hence is convex in x. Thus, for fixed $u \in \widehat{U}$, $f^i(x,u)$ is convex in x.

Moreover, the reformulation of the nonconvex QP over an ellipsoid into a convex QP over the same ellipsoid via the relation between $u^\top Q_x^i u + 2(r_x^i)^\top u + s_x^i$ and $f^i(x,u)$ in (28) follows from [21, Theorem 2.7]. □

Lemma B.1 implies that $\sup_{u \in \widehat{U}} f^i(x,u) \le 0$ is an alternate representation of our robust quadratic constraint. We next state the convergence rate in our framework for the associated feasibility problem. For this, we define the quantities

$$\sigma^2 := \max_{i \in [m]} \sum_{k=1}^K \|P_k^i\|_{\text{Fro}}^2, \quad \chi := \max_{i \in [m]} \max_{k \in [K]} \|P_k^i\|_{\text{Spec}}, \quad \text{and}$$
$$\rho := \max_{i \in [m]} \|A_i\|_{\text{Spec}}, \quad \beta := \max_{i \in [m]} \|b_i\|_2, \quad \mathcal{C} := \max_{i \in [m]} |c_i|. \tag{29}$$

Note that $\chi \le \sigma$. Furthermore, [5, Lemma 7] proves that $\|Q_x^i\|_{\text{Fro}} \le \sigma^2$ and $\|r_x^i\|_2 \le \sigma\rho$ holds for all x such that $\|x\|_2 \le 1$.

Corollary B.1. *Let our domain be given by $X = \{x \in \mathbb{R}^n : \|x\|_2 \le 1\}$. The customization of our OFO-based approach to the problem (27) ensures that within $O\left((\rho + \sqrt{K}\sigma)^2 + \sigma^2 + \beta + \mathcal{C}\right)^2 \log(m)\epsilon^{-2}$ iterations, we obtain a robust feasibility/infeasibility certificate. Moreover, each iteration in our framework relies on a first-order update where the most expensive operation in the case of (27) is computing $\lambda_{\max}(Q_x^i)$, which can be done in linear time.*

Proof. In order to apply OFO-based approach, we need to customize our proximal setup. Given that the sets X and $\widehat{U}$ are Euclidean balls, we set the proximal setup for generating the iterates $\{x_t, u_t^i\}_{t=1}^T$ to be the standard Euclidean d.g.f. with $\|\cdot\|_2$-norm. Then $\Omega_X = \Omega_{\widehat{U}} = \frac{1}{2}$.

Then we must bound the magnitude of the gradients measured by the $\|\cdot\|_2$-norm. Note that for any $i \in [m]$, the gradients of f^i are given by

$$\nabla_u f^i(x,u) = 2\left(Q_x^i - \lambda_{\max}(Q_x^i) I_K\right) u + 2 r_x^i$$
$$\nabla_x f^i(x,u) = 2 \left(A_i + \sum_{k=1}^K P_k^i \, u^{(k)} \right)^\top \left(A_i + \sum_{k=1}^K P_k^i \, u^{(k)} \right) x$$
$$+ 2\left(1 - \|u\|_2^2\right) \left(\sum_{k=1}^K P_k^i \, v^{(k)} \right)^\top \left(\sum_{k=1}^K P_k^i \, v^{(k)} \right) x - b_i,$$

where $v \in \widehat{U}$ is an eigenvector of Q_x^i corresponding to $\lambda_{\max}(Q_x^i)$.

Let us fix an $i \in [m]$. We first bound $\|\nabla_u f^i(x,u)\|_2$ for any $u \in \widehat{U}$ as follows:

$$\begin{aligned}\|\nabla_u f^i(x,u)\|_2 &= 2\left\|\left(Q_x^i - \lambda_{\max}(Q_x^i) I_K\right) u + r_x^i\right\|_2 \\ &\le 2\left(\left\|\left(Q_x^i - \lambda_{\max}(Q_x^i) I_K\right) u\right\|_2 + \left\|r_x^i\right\|_2\right) \\ &\le 2\lambda_{\max}(Q_x^i)\|u\|_2 + 2\sigma\rho \\ &\le 2(\sigma^2 + \sigma\rho),\end{aligned}$$

where the second inequality follows from $\left\|Q_x^i - \lambda_{\max}(Q_x^i) I_K\right\|_{\text{Spec}} \le \lambda_{\max}(Q_x^i)$ and $\|r_x^i\|_2 \le \sigma\rho$ which is implied by [5, Lemma 7], and the last inequality follows from the facts that $u \in \widehat{U}$, the definitions given in (29), and $\lambda_{\max}(Q_x^i) = \|\mathcal{P}_x^i\|_{\text{Spec}}^2 \le \|\mathcal{P}_x^i\|_{\text{Fro}}^2 \le \sum_{k=1}^K \|P_k^i\|_{\text{Fro}}^2 \le \sigma^2$ for any $x \in X$. Therefore, we deduce from Theorem 2.1 with uniform weights $\theta_t = \frac{1}{T}$ that the rate of convergence for bounding the weighted regret associated with constraint $i \in [m]$ using the online subgradient/mirror descent algorithm of [20, Theorem 1] is

$$\sup_{u \in U} \frac{1}{T}\sum_{t=1}^T f^i(x_t,u) - \frac{1}{T}\sum_{t=1}^T f^i(x_t,u_t) \le \frac{2(\sigma^2+\sigma\rho)}{\sqrt{T}}.$$

This implies that $r_u(\epsilon) = O((\sigma^2+\sigma\rho)^2\epsilon^{-2})$.

Before bounding $\|\nabla_x f^i(x,u)\|_2$, we observe that for any $u \in \widehat{U}$

$$\left\|\sum_{k=1}^K P_k^i\, u^{(k)}\right\|_{\text{Spec}} \le \sum_{k=1}^K \|P_k^i\|_{\text{Spec}}\, |u^{(k)}| \le \sqrt{K} \max_{k\in[K]} \|P_k^i\|_{\text{Spec}} \le \sqrt{K}\chi,$$

where the second inequality holds because $\|u\|_1 \le \sqrt{K}\|u\|_2 \le \sqrt{K}$ holds for all $u \in \widehat{U}$. Then for any $x \in X$, $u \in \widehat{U}$, and eigenvector $v \in \widehat{U}$, we have

$$\begin{aligned}\|\nabla_x f^i(x,u)\|_2 &\le 2\left\|A_i + \sum_{k=1}^K P_k^i\, u^{(k)}\right\|_{\text{Spec}}^2 \|x\|_2 + 2\left(1 - \|u\|_2^2\right)\left\|\sum_{k=1}^K P_k^i\, v^{(k)}\right\|_{\text{Spec}}^2 \|x\|_2 + \|b_i\|_2 \\ &\le 2(\rho + \sqrt{K}\chi)^2 + 2K\chi^2 + \beta \\ &\le 4(\rho + \sqrt{K}\chi)^2 + \beta \\ &\le 4(\rho + \sqrt{K}\sigma)^2 + \beta.\end{aligned}$$

We next bound the online SP gap of the functions $\phi_t(x,y) = \sum_{i=1}^m y^{(i)} f^i(x,u_t^i)$, where $y \in Y = \Delta_m$. In this situation, the proximal setup for Y is the negative entropy with $\|\cdot\|_1$-norm, thus $\Omega_Y = \log(m)$. We must bound the dual norm of $\nabla_y \phi_t(x,y)$, that is,

$$\begin{aligned}\|\nabla_y \phi_t(x,y)\|_\infty &= \left\|\left\{f^i(x,u_t^i)\right\}_{i\in[m]}\right\|_\infty = \max_{i\in[m]} |f^i(x,u_t^i)| \\ &= \max_{i\in[m]}\left\{\left|\left\|\left(A_i + \sum_{k=1}^K P_k^i\, u_t^{i,(k)}\right) x\right\|_2^2 - b_i^\top x - c_i + \lambda_{\max}(Q_x^i)\left(1 - \|u_t^i\|_2^2\right)\right|\right\} \\ &\le (\rho + \sqrt{K}\chi)^2 + \sigma^2 + \beta + \max_{i\in[m]} |c_i| \\ &\le (\rho + \sqrt{K}\sigma)^2 + \sigma^2 + \beta + \mathcal{C}.\end{aligned}$$

Also, because $y \in \Delta_m$, we have

$$\|\nabla_x \phi_t(x,y)\|_2 = \left\| \sum_{i=1}^m y^{(i)} \nabla_x f^i(x,u) \right\|_2 \leq \sum_{i=1}^m y^{(i)} \|\nabla_x f^i(x,u)\|_2 \leq 4(\rho + \sqrt{K}\sigma)^2 + \beta.$$

Taking into account the proximal setup for the joint space $X \times Y$ and the resulting constants (see [30, 25, 26] for details), we have $\Omega_{X\times Y} = 1$ and $\|[\nabla_x \phi_t(x,y); -\nabla_y \phi_t(x,y)]\|_{X\times Y} \leq L_{X\times Y}$ where

$$\begin{aligned} L_{X\times Y} &= \left((\rho + \sqrt{K}\sigma)^2 + \sigma^2 + \beta + \mathcal{C}\right)\sqrt{\log(m)} + \left(4(\rho + \sqrt{K}\sigma)^2 + \beta\right)\sqrt{1/2} \\ &\leq \left(4(\rho + \sqrt{K}\sigma)^2 + \sigma^2 + \beta + \mathcal{C}\right)\sqrt{\log(m)}. \end{aligned}$$

Then from Theorem 2.2 the rate of convergence for the online SP problem using the online subgradient/mirror descent algorithm of [20, Theorem 3] is

$$\sup_{y\in Y} \frac{1}{T}\sum_{t=1}^T \phi_t(x_t, y) - \inf_{x\in X} \frac{1}{T}\sum_{t=1}^T \phi_t(x, y_t) \leq \left(4(\rho + \sqrt{K}\sigma)^2 + \sigma^2 + \beta + \mathcal{C}\right)\sqrt{\frac{2\log(m)}{T}}.$$

This now implies that $r_x(\epsilon) = O\left((\rho + \sqrt{K}\sigma)^2 + \sigma^2 + \beta + \mathcal{C}\right)^2 \log(m)\epsilon^{-2}$. Therefore, the number of iterations required for our OFO-based approach to obtain a robust feasibility/infeasibility certificate is $T = \max\{r_x(\epsilon), r_u(\epsilon)\}$, which gives us the result.

Note that each iteration of our approach requires a first-order update that is composed of computing the gradients $\nabla_x f^i(x,u)$ and $\nabla_u f^i(x,u)$ and prox computations. Because our domains involve only direct products of Euclidean balls and simplices, they admit efficient prox computations which take $O(Km + mn)$ time. In order to evaluate the gradients $\nabla_x f^i(x,u)$ and $\nabla_u f^i(x,u)$, in addition to the elementary matrix vector operations, we need to compute $\lambda_{\max}(Q_x^i)$ which is the most expensive operation in our first-order update. Fortunately, the maximum eigenvalue of a matrix can be computed in linear time (see [17, Chapter 8.2]). □

In the case of robust QP feasibility problem (27), [5, Corollary 3] states that with probability $1-\delta$, their framework returns robust feasibility/infeasibility certificates in at most $O\left(K^2\sigma^2(\rho^2+\sigma^2)\log(m/\delta)\epsilon^{-2}\right)$ calls (iterations) to their oracle. In each call to their oracle, a nominal feasibility problem is solved to the accuracy $\epsilon/2$. In comparison we deduce from Corollary B.1 that our framework requires comparable number of iterations as the approach of Ben-Tal et al. [5]. Even so, there are a number of reasons that considerably favor our approach. First, our approach is deterministic as opposed to the high $1-\delta$ probability guarantee of [5] which requires using an adaptation of the follow-the-perturbed-leader type OCO. Second, each iteration of their approach requires solving a nominal feasibility problem for solution oracle as well as solving TRSs for the computation of u_t. In contrast to this, in each iteration we carry out mainly elementary operations such as matrix vector multiplications and our most computationally expensive operation is the maximum eigenvalue computations $\lambda_{\max}(Q_x^i)$. While there are established algorithms to solve the TRS, TRS is inherently more complicated than finding the maximum eigenvalue of a positive semidefinite matrix. Moreover, [5] suffers from the additional computational cost of their solution oracle which solves the nominal feasibility problem. Hence, our approach, while requiring a comparable number of iterations, reduces the cost per iteration remarkably.